

VMware Cloud Foundation 5.2

Table of Contents

Release Notes	26
VMware Cloud Foundation 5.2.1 Release Notes	26
Introduction	26
What's New	26
Available Languages	27
Important VMware Cloud Foundation 5.2.1 API Changes	27
Deprecation Notices	28
VMware Cloud Foundation Bill of Materials (BOM)	29
Supported Hardware	29
Documentation	29
Browser Compatibility and Screen Resolutions	29
Installation and Upgrade Information	30
VMware Cloud Foundation 5.2.1.1 Release Information	30
Resolved Issues	31
Known Issues	31
VMware Cloud Foundation Known Issues	31
Upgrade Known Issues	32
Bring-up Known Issues	35
SDDC Manager Known Issues	35
Workload Domain Known Issues	36
VMware Cloud Foundation 5.2.1 on Dell VxRail Release Notes	41
Introduction	41
What's New	41
Available Languages	41
Deprecation Notices	42
VMware Cloud Foundation Bill of Materials (BOM)	42
Documentation	43
Installation and Upgrade Information	43
VMware Cloud Foundation 5.2.1.1 Release Information	43
Resolved Issues	44
Known Issues	44
VMware Cloud Foundation 5.2 Release Notes	46
Introduction	46
What's New	46
Available Languages	47
Deprecation Notices	48
VMware Cloud Foundation Bill of Materials (BOM)	48

Supported Hardware	49
Documentation	49
Browser Compatibility and Screen Resolutions	49
Installation and Upgrade Information.....	49
Resolved Issues	50
Known Issues	50
VMware Cloud Foundation Known Issues	50
Upgrade Known Issues	51
Bring-up Known Issues.....	54
SDDC Manager Known Issues.....	54
Workload Domain Known Issues.....	55
VMware Cloud Foundation 5.2 on Dell VxRail Release Notes	59
Introduction.....	60
What's New.....	60
Available Languages	60
Deprecation Notices	61
VMware Cloud Foundation Bill of Materials (BOM).....	61
Documentation	61
Installation and Upgrade Information.....	62
Resolved Issues	62
Known Issues	62
Async Patch Tool Release Notes	64
Introduction.....	65
What's New.....	65
Resolved Issues	65
Known Issues	65
Design Guide	69
Intended Audience	69
Before You Apply This Guidance	69
Design Elements	69
VMware Cloud Foundation Deployment Options in This Design.....	69
vCenter Single Sign-On Options in This Design.....	70
VMware Cloud Foundation Design Blueprints.....	70
Information about Environment Configurations that Can be Converted or Imported into VMware Cloud Foundation.....	70
<i>VMware Cloud Foundation Glossary.....</i>	<i>70</i>
<i>VMware Cloud Foundation Concepts.....</i>	<i>70</i>
Architecture Models and Workload Domain Types in VMware Cloud Foundation	70
Architecture Models	70
Workload Domain Types	71

Workload Domain Cluster to Rack Mapping in VMware Cloud Foundation	75
Networking Models in VMware Cloud Foundation	78
External Services Design for VMware Cloud Foundation.....	79
Supported Storage Types for VMware Cloud Foundation	79
vSphere Design for VMware Cloud Foundation.....	80
NSX Design for VMware Cloud Foundation	80
Logical Design for NSX for VMware Cloud Foundation	82
Single Instance - Single Availability Zone	83
Single Instance - Multiple Availability Zones	84
Multiple Instances - Single Availability Zone	85
Multiple Instances - Multiple Availability Zones	86
Routing Design for VMware Cloud Foundation	87
Routing Options for VMware Cloud Foundation	87
BGP Routing Design for VMware Cloud Foundation	89
BGP Routing Design Requirements and Recommendations for VMware Cloud Foundation.....	94
Lifecycle Management Design for VMware Cloud Foundation	102
VMware Cloud Foundation Lifecycle Management Requirements	103
Logging and Monitoring Design for VMware Cloud Foundation.....	104
VMware Cloud Foundation Topology Design Blueprints.....	104
Topology Design Blueprint One: Single Instance - Single Availability Zone	105
Design Choices for Design Blueprint One	105
Design Elements for Design Blueprint One	106
Topology Design Blueprint Two: Consolidated Single Instance - Single Availability Zone	107
Design Choices for Design Blueprint Two	108
Design Elements for Design Blueprint Two.....	109
Topology Design Blueprint Three: Single Instance - Multiple Availability Zones	110
Design Choices for Design Blueprint Three	111
Design Elements for Design Blueprint Three	112
Topology Design Blueprint Four: Multiple Instance - Single Availability Zone	115
Design Choices for Design Blueprint Four	116
Design Elements for Design Blueprint Four	117
Topology Design Blueprint Five: Multiple Instance - Multiple Availability Zones	120
Design Choices for Design Blueprint Five	121
Design Elements for Design Blueprint Five	122
VMware Cloud Foundation vSphere Cluster Design Patterns	125
vSphere Cluster Design Pattern One: Multi-Rack Compute VI Workload Domain Cluster	126
Design Choices for vSphere Cluster Design Pattern One	126
Design Elements for vSphere Cluster Design Pattern One	127
VMware Cloud Foundation NSX Edge Cluster Design Patterns	128
NSX Edge Cluster Design Pattern One: Dedicated Edge Scale and Performance	128

Design Choices for Edge Cluster Design Pattern One	129
Design Elements for NSX Edge Cluster Design Pattern One	130
NSX Edge Cluster Design Pattern Two: Multi-Rack Edge Availability	131
Design Choices for NSX Edge Cluster Design Pattern Two	132
Design Elements for NSX Edge Cluster Design Pattern Two	133
Appendix: Design Elements for VMware Cloud Foundation	134
Architecture Design Elements for VMware Cloud Foundation	134
Workload Domain Design Elements for VMware Cloud Foundation	134
External Services Design Elements for VMware Cloud Foundation	135
Physical Network Design Elements for VMware Cloud Foundation	136
vSAN Design Elements for VMware Cloud Foundation	140
ESXi Design Elements for VMware Cloud Foundation	147
vCenter Server Design Elements for VMware Cloud Foundation	150
vCenter Server Design Elements	150
vCenter Single Sign-On Design Elements	152
vSphere Cluster Design Elements for VMware Cloud Foundation	153
vSphere Networking Design Elements for VMware Cloud Foundation	157
NSX Design Elements for VMware Cloud Foundation	160
NSX Manager Design Elements	160
NSX Global Manager Design Elements	162
NSX Edge Design Elements	164
BGP Routing Design Elements for VMware Cloud Foundation	168
Overlay Design Elements for VMware Cloud Foundation	175
Application Virtual Network Design Elements for VMware Cloud Foundation	177
Load Balancing Design Elements for VMware Cloud Foundation	179
SDDC Manager Design Elements for VMware Cloud Foundation	181
VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation	182
Workspace ONE Access Design Elements for VMware Cloud Foundation	186
Lifecycle Management Design Elements for VMware Cloud Foundation	191
Information Security Design Elements for VMware Cloud Foundation	193
Planning and Preparation Workbook	195
Download the Planning and Preparation Workbook	195
Send Us Your Feedback on the Planning and Preparation Workbook	195
Getting Started with VMware Cloud Foundation	196
Intended Audience	196
Related VMware Cloud Foundation Publications	196
VMware Cloud Foundation Glossary	196
VMware Cloud Foundation Overview	196
VMware Cloud Foundation Components	197
VMware Cloud Foundation Features	199

VMware Cloud Foundation Glossary	200
Deployment Guide.....	203
About the VMware Cloud Foundation Deployment Guide	203
Intended Audience	203
Related Publications	203
VMware Cloud Foundation Glossary	203
Preparing your Environment for VMware Cloud Foundation	203
Deploying VMware Cloud Foundation	203
Deploy VMware Cloud Builder Appliance	204
Prepare ESXi Hosts for VMware Cloud Foundation.....	206
Create a Custom ISO Image for ESXi.....	206
Install ESXi Interactively and Configure Hosts for VMware Cloud Foundation	208
Regenerate the Self-Signed Certificate on All Hosts	211
Configure ESXi Hosts with Signed Certificates.....	212
Deploy the Management Domain Using VMware Cloud Builder	213
About the Deployment Parameter Workbook.....	214
Deploy the Management Domain Using ESXi Hosts with External Certificates	225
Troubleshooting VMware Cloud Foundation Deployment.....	227
VMware Cloud Builder Log Files	227
Using the SoS Utility on VMware Cloud Builder.....	227
SoS Utility Help Options	227
SoS Utility Generic Options	228
SoS Utility Log File Options	228
SoS Utility JSON Generator Options	229
SoS Utility Health Check Options.....	229
Sample Output	230
VMware Cloud Foundation Glossary	200
Administration Guide.....	234
About the VMware Cloud Foundation Administration Guide	234
Intended Audience	234
Related Publications	234
Administering VMware Cloud Foundation	234
VMware Software Components Deployed by VMware Cloud Foundation	235
Web Interfaces Used to Administer VMware Cloud Foundation	235
Getting Started with SDDC Manager.....	235
Log in to the SDDC Manager User Interface.....	236
Guided SDDC Manager Onboarding.....	236
Tour of the SDDC Manager User Interface	237
Navigation Bar.....	237
Log out of the SDDC Manager User Interface	239

Configure the Customer Experience Improvement Program Settings for VMware Cloud Foundation.....	239
Managing Certificates in VMware Cloud Foundation	240
View Certificate Information	240
Configure VMware Cloud Foundation to Use Microsoft CA-Signed Certificates.....	241
Prepare Your Microsoft Certificate Authority to Allow SDDC Manager to Manage Certificates	241
Configure a Microsoft Certificate Authority in SDDC Manager	245
Install Microsoft CA-Signed Certificates using SDDC Manager.....	246
Configure VMware Cloud Foundation to Use OpenSSL CA-Signed Certificates.....	247
Configure OpenSSL-signed Certificates in SDDC Manager.....	247
Install OpenSSL-signed Certificates using SDDC Manager	248
Install Third-Party CA-Signed Certificates Using Server Certificate and Certificate Authority Files	250
Install Third-Party CA-Signed Certificates in VMware Cloud Foundation Using a Certificate Bundle.....	252
Add a Trusted Certificate to the SDDC Manager Trust Store	254
Remove Old or Unused Certificates from SDDC Manager	255
Managing License Keys in VMware Cloud Foundation.....	255
Add a Component License Key in the SDDC Manager UI	256
Edit a Component License Key Description in the SDDC Manager UI	256
Delete a Component License Key in the SDDC Manager UI	257
Update Component License Keys for Workload Domain Components.....	257
Prepare ESXi Hosts for VMware Cloud Foundation	206
Create a Custom ISO Image for ESXi	206
Create a Custom ESXi ISO Image Using VMware PowerCLI	207
Create a Custom ESXi ISO Image Using vSphere Lifecycle Manager.....	208
Install ESXi Interactively and Configure Hosts for VMware Cloud Foundation	208
Install ESXi on VMware Cloud Foundation Hosts Using the ISO	208
Configure the Network on VMware Cloud Foundation Hosts	209
Configure the Virtual Machine Network Port Group on VMware Cloud Foundation Hosts.....	210
Configure NTP on VMware Cloud Foundation Hosts.....	210
Regenerate the Self-Signed Certificate on All Hosts	211
Configure ESXi Hosts with Signed Certificates	212
Managing ESXi Hosts in VMware Cloud Foundation	265
Network Pool Management	265
Size a Network Pool.....	266
View Network Pool Details	267
Create a Network Pool	267
Add or Remove a Network Pool IP Address Range.....	269
Rename a Network Pool	269
Delete a Network Pool	270
View Host Inventory.....	270
Commission Hosts.....	271

Decommission Hosts	274
ESXi Lockdown Mode	275
Managing vSphere Lifecycle Manager Images in VMware Cloud Foundation.....	275
vSphere Lifecycle Manager Image Components	276
vSphere Lifecycle Manager Images in VMware Cloud Foundation.....	276
vSphere Lifecycle Manager Image Workflow	277
vSphere Lifecycle Manager Options for Workload Domains	277
Create a vSphere Lifecycle Manager Image	277
Export a vSphere Lifecycle Manager Image	279
Creating a vSphere Lifecycle Manager Image in VMware Cloud Foundation.....	280
Extract a vSphere Lifecycle Manager Image	280
Import a vSphere Lifecycle Manager Image	280
Firmware Updates	282
View vSphere Lifecycle Manager Images	282
Managing Storage in VMware Cloud Foundation	282
Principal Storage	282
Supplemental Storage	283
vSAN Storage with VMware Cloud Foundation.....	283
Prerequisites for vSAN Storage	283
Procedures for vSAN Storage.....	284
vSAN Original Storage Architecture (OSA).....	284
vSAN Express Storage Architecture (ESA).....	284
vSAN Max	284
vSAN Compute Clusters	284
NFS Storage with VMware Cloud Foundation.....	285
Prerequisites for NFS Storage	285
Procedures for NFS Storage.....	285
Fibre Channel Storage with VMware Cloud Foundation	286
Prerequisites for FC Storage.....	286
Procedures for FC Storage	286
HCI Mesh with VMware Cloud Foundation	286
HCI Mesh Compute-only Clusters.....	287
vVols Storage with VMware Cloud Foundation	287
Prerequisites for vVols Storage.....	288
Procedures for vVols Storage	288
Add a VASA Provider	288
View a VASA Provider	289
Edit a VASA Provider	289
Delete a VASA Provider	289
Converting or Importing Existing vSphere Environments into VMware Cloud Foundation	289

Supported Scenarios for Converting or Importing vSphere Environments to VMware Cloud Foundation	289
Supported Scenarios.....	290
Glossary of Terms for Converting/Importing vSphere Environments	290
Considerations Before Converting or Importing Existing vSphere Environments into VMware Cloud Foundation.	290
Considerations by Domain Type	290
Configuration Support by Domain Type	294
Network and Compute Requirements for Converting or Importing Existing vSphere Environments	296
Network and Compute Requirements.....	296
Download Software for Converting or Importing Existing vSphere Environments	296
VCF Import Tool Options and Parameters.....	298
Convert a vSphere Environment to a Management Domain or Import a vSphere Environment as a VI Workload Domain in VMware Cloud Foundation.....	299
Copy the VCF Import Tool to the Target vCenter Appliance.....	300
Run a Precheck on the Target vCenter Before Conversion	301
Deploy the SDDC Manager Appliance on the Target vCenter	302
Generate an NSX Deployment Specification for Converting or Importing Existing vSphere Environments	303
Upload the Required Software to the SDDC Manager Appliance.....	304
Run a Detailed Check on the Target vCenter Before Conversion or Import	305
Convert or Import the vSphere Environment into the SDDC Manager Inventory	305
Add Licenses for Converted or Imported Workload Domains in SDDC Manager.....	306
Validate a Converted Management Domain or Imported VI Workload Domain	306
VCF Import Tool Troubleshooting	307
General Troubleshooting for the VCF Import Tool.....	307
VCF Import Tool Guardrail Troubleshooting.....	307
Managing Workload Domains in VMware Cloud Foundation	311
About VI Workload Domains	312
Prerequisites for a Workload Domain	312
Deploy NSX Manager for Workload Domains	314
Delete a VI Workload Domain	317
View Workload Domain Details	317
Expand a Workload Domain.....	319
Add a Host to a vSphere Cluster Using the SDDC Manager UI	319
Add a vSphere Cluster to a Workload Domain Using the SDDC Manager UI	321
Shrink a Workload Domain	326
Remove a Host from a vSphere Cluster in a Workload Domain	326
Delete a vSphere Cluster from a Workload Domain	327
Rename a Workload Domain	327
vSphere Cluster Management.....	328
View vSphere Cluster Details.....	328
Rename a Cluster in the SDDC Manager UI	328

Mount a Remote vSAN Datastore	329
Unmount a Remote vSAN Datastore	330
Tag Management	330
Tag a Workload Domain	330
Tag a Cluster	331
Tag a Host	332
Manage Workload Domain Configuration Drift Between vCenter Server and SDDC Manager	333
Managing NSX Edge Clusters in VMware Cloud Foundation	334
Prerequisites for an NSX Edge Cluster	335
Deploy an NSX Edge Cluster	335
Add Edge Nodes to an NSX Edge Cluster	341
Remove Edge Nodes from an NSX Edge Cluster	345
Managing Avi Load Balancer in VMware Cloud Foundation	346
Limitations of Avi Load Balancer in VMware Cloud Foundation	347
Deploy Avi Load Balancer for a Workload Domain	347
Remove Avi Load Balancer from a Workload Domain	350
Deploying Application Virtual Networks in VMware Cloud Foundation	351
Overlay-Backed NSX Segments	351
VLAN-Backed NSX Segments	351
VMware Aria Suite Components and NSX Segments	351
Identity and Access Management for VMware Cloud Foundation	352
Deploy Overlay-Backed NSX Segments	352
Deploy VLAN-Backed NSX Segments	353
VMware Cloud Foundation with VMware Tanzu	354
Enable Workload Management	355
View Workload Management Cluster Details	356
Update Workload Management License	356
VMware Aria Suite Lifecycle in VMware Cloud Foundation mode	356
VMware Aria Suite Lifecycle Implementation	357
Deploy VMware Aria Suite Lifecycle	358
Replace the Certificate of the VMware Aria Suite Lifecycle Instance	359
Configure Data Center and vCenter Server in VMware Aria Suite Lifecycle	359
Workspace ONE Access Implementation	360
Import the Workspace ONE Access Certificate to VMware Aria Suite Lifecycle	361
Add Workspace ONE Access Passwords to VMware Aria Suite Lifecycle	362
Deploy a Standard Workspace ONE Access Instance Using VMware Aria Suite Lifecycle	363
Deploy Clustered Workspace ONE Access Instance Using VMware Aria Suite Lifecycle	365
Configure an Anti-Affinity Rule and a Virtual Machine Group for a Clustered Workspace ONE Access Instance	367
Configure NTP on Workspace ONE Access	368

Configure the Domain and Domain Search Parameters on Workspace ONE Access	368
Configure an Identity Source for Workspace ONE Access	369
Add the Clustered Workspace ONE Access Cluster Nodes as Identity Provider Connectors	370
Assign Roles to Active Directory Groups for Workspace ONE Access	371
Assign Roles to Active Directory Groups for VMware Aria Suite Lifecycle	371
Working with NSX Federation in VMware Cloud Foundation	371
NSX Federation Key Concepts.....	372
NSX Federation Systems: Global Manager and Local Manager	372
NSX Federation Span: Local and Cross-Instance	372
NSX Federation Tunnel Endpoints.....	372
NSX Federation Tier Gateways.....	372
Configuring NSX Federation in VMware Cloud Foundation	373
Create Global Manager Clusters for VMware Cloud Foundation.....	374
Prepare Local Manager for NSX Federation in VMware Cloud Foundation	377
Enable NSX Federation in VMware Cloud Foundation.....	377
Stretch Segments between VMware Cloud Foundation Instances.....	379
Set Standby Global Manager	382
Replacing Global Manager Cluster Certificates in VMware Cloud Foundation	382
Import a CA-Signed Certificate to the Global Manager Cluster	383
Replace the Certificate for the First Global Manager Node	383
Replace Certificates and Virtual IP for the Remaining Global Manager Nodes	384
Update Local Manager Certificate Thumbprint in Global Manager Cluster	386
Password Management for NSX Global Manager Cluster in VMware Cloud Foundation.....	386
Backup and Restore of NSX Global Manager Cluster in VMware Cloud Foundation.....	386
Configure NSX Global Manager Cluster Backups	387
Restore an NSX Global Manager Cluster Backup	387
Managing Installation and Upgrade Bundles in VMware Cloud Foundation	388
Bundle Types	388
Downloading Install Bundles for VMware Cloud Foundation	389
Connect SDDC Manager to a Software Depot for Downloading Bundles	389
Download an Install Bundle from SDDC Manager	390
Configure a Proxy Server for Downloading Bundles.....	391
Download an Install Bundle Using the Bundle Transfer Utility	391
View Bundle Download History	394
Stretching vSAN Clusters in VMware Cloud Foundation.....	394
About Availability Zones and Regions	395
Availability Zones	395
Regions	395
Stretched Cluster Requirements	395
VLANs and Subnets for Multiple Available Zones.....	395

Networking for Multiple Availability Zones	396
Deploy and Configure vSAN Witness Host	397
Deploy vSAN Witness Host.....	397
Register vSAN Witness Host	398
Configure NTP on the Witness Host	399
Configure the VMkernel Adapters on the vSAN Witness Host.....	399
Stretch a vSAN Cluster in VMware Cloud Foundation	400
NSX Configuration for Availability Zone 2.....	416
Configure IP Prefixes in the Tier-0 Gateway for Availability Zone 2.....	416
Configure Route Maps in the Tier-0 Gateway for Availability Zone 2.....	417
Configure BGP in the Tier-0 Gateway for Availability Zone 2.....	418
Expand a Stretched Cluster in VMware Cloud Foundation	419
Unstretch a Cluster	422
Replace a Failed Host in a Stretched Cluster	423
Change the vSAN Witness Host in a Stretched Cluster	424
Monitoring Capabilities in the VMware Cloud Foundation System	425
Viewing Tasks and Task Details.....	425
Viewing and Filtering Task Details	425
Managing Tasks and Subtask Details	426
Resizing the Task Panel	426
API Activity Logging.....	426
Activity Log Structure	426
Activity Log Example.....	426
Activity Logs Retention Policy	427
Log Analysis	427
Updating VMware Cloud Foundation DNS and NTP Servers.....	427
Update DNS Server Configuration	427
Update NTP Server Configuration.....	428
Supportability and Serviceability (SoS) Utility	429
SoS Utility Options.....	429
SoS Utility Help Options	429
SoS Utility Generic Options	430
SoS Utility VMware Cloud Foundation Summary Options	431
SoS Utility Fix-It-Up Options	431
SoS Utility Health Check Options.....	433
Example Health Check Commands:	434
Collect Logs for Your VMware Cloud Foundation System.....	434
Component Log Files Collected by the SoS Utility	437
Replacing Host Components in VMware Cloud Foundation	438
Avoiding Unintentional Downtime.....	439

Replacing Components of a Host Running in Degraded Mode.....	439
Replace Components of an Assigned Host Running in Degraded Mode	439
Replace Components of an Unassigned Host Running in Degraded Mode	440
Replace a Dead Host	440
Replace Boot Disk on a Host.....	441
Managing Users and Groups in VMware Cloud Foundation	441
Configuring the Identity Provider for VMware Cloud Foundation	442
Add Active Directory over LDAP or OpenLDAP as an Identity Source for VMware Cloud Foundation	442
Configure Microsoft ADFS as the Identity Provider in the SDDC Manager UI.....	444
Add a User or Group to VMware Cloud Foundation.....	446
Remove a User or Group	447
Create a Local Account	447
Create an Automation Account.....	449
Managing Passwords in VMware Cloud Foundation.....	453
Password Expiration Notifications	453
Rotate Passwords	454
Manually Update Passwords	456
Remediate Passwords	456
Look Up Account Credentials	457
Updating SDDC Manager Passwords	458
Update SDDC Manager Root and Super User Passwords.....	458
Update SDDC Manager Local Account Password	458
Update Expired SDDC Manager Root Password	459
Backup and Restore of VMware Cloud Foundation	459
Reconfigure SFTP Backups for SDDC Manager and NSX Manager	460
File-Based Backups for SDDC Manager and vCenter Server.....	461
Back Up SDDC Manager	462
Configure a Backup Schedule for vCenter Server	463
Manually Back Up vCenter Server	463
Export the Configuration of the vSphere Distributed Switches	464
File-Based Restore for SDDC Manager, vCenter Server, and NSX.....	465
Restore SDDC Manager	465
Restore vCenter Server	468
Restore the Configuration of a vSphere Distributed Switch	474
Restore an NSX Manager Cluster Node.....	474
Restoring NSX Edge Cluster Nodes	483
Image-Based Backup and Restore of VMware Cloud Foundation.....	490
VMware Cloud Foundation Glossary	200
Operations Guide	493
Intended Audience	493

PowerShell Modules for VMware Cloud Foundation Operations	493
Related VMware Cloud Foundation Publications	494
VMware Cloud Foundation Glossary	494
Shutdown and Startup of VMware Cloud Foundation	494
Shutting Down and Starting Up VMware Cloud Foundation by Using PowerShell	494
Shutting Down VMware Cloud Foundation	495
Starting Up VMware Cloud Foundation	495
Password Policy Configuration for VMware Cloud Foundation	495
Password Policy Configuration and Password Management	495
Manual and Automated Password Policy Configuration.....	496
Approaches to Password Policy Configuration	496
Prerequisites.....	497
Lifecycle Management Guide	499
Upgrading VMware Cloud Foundation to 5.2.x	499
SDDC Manager Functionality During an Upgrade to VMware Cloud Foundation 5.2.....	499
Upgrade States and Terminology.....	499
vSphere UI Client Plug-ins	502
Monitor VMware Cloud Foundation Updates	502
View VMware Cloud Foundation Update History	504
Access VMware Cloud Foundation Upgrade Log Files	504
Downloading VMware Cloud Foundation Upgrade Bundles	505
Online and Offline Downloads	505
Other Bundle Types.....	505
Connect SDDC Manager to a Software Depot for Downloading Bundles.....	389
Download Bundles Using SDDC Manager.....	507
Configure a Proxy Server for Downloading VMware Cloud Foundation Bundles.....	507
Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles	508
Offline Download of Independent SDDC Manager Bundles.....	511
Offline Download of Async Patch Bundles	513
Offline Download of Flexible BOM Upgrade Bundles.....	515
HCL Offline Download for VMware Cloud Foundation	517
Download Bundles to an Offline Depot.....	519
VMware Cloud Foundation Upgrade Prerequisites	520
VMware Cloud Foundation 5.2.x Upgrade Overview	521
VMware Cloud Foundation Upgrade Preparation	521
Management Domain Upgrade	521
VI Workload Domain Upgrade.....	524
Upgrade the Management Domain to VMware Cloud Foundation 5.2.x	526
Perform Update Precheck - Versions Prior to SDDC Manager 5.0	527
Perform Update Precheck in SDDC Manager	529

Apply the VMware Cloud Foundation 5.2.x Upgrade Bundle	533
Apply VMware Cloud Foundation Configuration Updates	534
Upgrade VMware Aria Suite Lifecycle and VMware Aria Suite Products for VMware Cloud Foundation	538
Upgrade NSX for VMware Cloud Foundation in a Federated Environment	538
Download NSX Global Manager Upgrade Bundle	538
Upgrade the Upgrade Coordinator for NSX Federation	539
Upgrade NSX Global Managers for VMware Cloud Foundation	539
Upgrade NSX for VMware Cloud Foundation 5.2.x	540
Upgrade vCenter Server for VMware Cloud Foundation 5.2.x	542
Upgrade ESXi for VMware Cloud Foundation 5.2.1	545
Upgrade ESXi with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation 5.2	546
Upgrade vSAN Witness Host for VMware Cloud Foundation	547
Skip Hosts During ESXi Update	548
Upgrade ESXi with Custom ISOs	549
Upgrade ESXi with VMware Cloud Foundation Stock ISO and Async Drivers	552
Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2	554
Firmware Updates	282
Update License Keys for a Workload Domain	558
Upgrade vSphere Distributed Switch versions	559
Upgrade vSAN on-disk format versions	559
Upgrade VI Workload Domains to VMware Cloud Foundation 5.2.x	560
Plan VI Workload Domain Upgrade	560
Perform Update Precheck in SDDC Manager	529
Upgrade NSX for VMware Cloud Foundation in a Federated Environment	538
Download NSX Global Manager Upgrade Bundle	538
Upgrade the Upgrade Coordinator for NSX Federation	539
Upgrade NSX Global Managers for VMware Cloud Foundation	539
Upgrade NSX for VMware Cloud Foundation 5.2.x	540
Upgrade vCenter Server for VMware Cloud Foundation 5.2.x	542
Upgrade ESXi for VMware Cloud Foundation 5.2.1	545
Upgrade ESXi with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation	571
Upgrade vSAN Witness Host for VMware Cloud Foundation	547
Skip Hosts During ESXi Update	548
Upgrade ESXi with Custom ISOs	549
Upgrade ESXi with VMware Cloud Foundation Stock ISO and Async Drivers	552
Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2	554
Firmware Updates	282
Update License Keys for a Workload Domain	558
Upgrade vSphere Distributed Switch versions	559
Upgrade vSAN on-disk format versions	559

Post Upgrade Steps for NFS-Based VI Workload Domains.....	583
Independent SDDC Manager Upgrade using the SDDC Manager UI	584
Flexible BOM Upgrade in VMware Cloud Foundation	584
Patching the Management and Workload Domains	587
Troubleshooting for Upgrading VMware Cloud Foundation	588
SDDC Manager Troubleshooting.....	588
On-demand pre-checks for vCenter bundle might fail	588
SDDC Manager bundle pre-check failure when upgrading to VMware Cloud Foundation 5.1	589
Extra RPM packages on SDDC Manager may cause upgrade failure	589
False warning for missing compatibility data in plan upgrade wizard	589
Updating licenses for a WLD shows insufficient license error	589
vCenter Troubleshooting	589
vCenter Server Upgrade Failed Due to Reuse of Temporary IP Address.....	590
Async Patch Tool.....	591
Async Patch Tool 1.2	591
Apply an Async Patch to VMware Cloud Foundation in Online Mode.....	592
Apply an Async Patch to VMware Cloud Foundation in Offline Mode	595
Upgrade an Async Patched Version of VMware Cloud Foundation in Online Mode.....	598
Upgrade an Async Patched Version of VMware Cloud Foundation in Offline Mode.....	600
VCF Async Patch Tool Options	603
Async Patch Tool Help Option	603
Customer Experience Improvement Program (CEIP) Option.....	604
List Async Patches Option.....	604
Download Patch Option (offline only)	605
Enable Patch Option.....	605
Precheck Option	606
Postcheck Option	607
Deactivate All Patches Option	608
Enable VCF Upgrade Option.....	608
Inventory Sync Option	609
VCF on Dell VxRail Guide	611
About VMware Cloud Foundation on Dell VxRail	611
Intended Audience	611
Related Publications.....	611
VMware Cloud Foundation on Dell VxRail.....	611
Prepare a VxRail Environment for Cloud Builder Appliance Deployment	612
Imaging the VxRail Management Nodes	612
VxRail First Run for the Management Cluster	612
Deploy VMware Cloud Builder Appliance.....	612
Deploy the Management Domain Using VMware Cloud Builder	213

About the Deployment Parameter Workbook	214
VxRail Prerequisites	616
Credentials Worksheet	214
Hosts and Networks Worksheet	216
Deploy Parameters Worksheet: Existing Infrastructure Details	220
Deploy Parameters Worksheet: VxRail Manager Details	622
Deployment Parameters Worksheet: License Keys	221
Deploy Parameters Worksheet: vSphere Infrastructure	221
Deploy Parameters Worksheet: VMware NSX	224
Deploy Parameters Worksheet: SDDC Manager	225
Troubleshooting VMware Cloud Foundation Deployment	227
Using the SoS Utility on VMware Cloud Builder	227
SoS Utility Help Options	227
SoS Utility Generic Options	228
SoS Utility Log File Options	228
SoS Utility JSON Generator Options	229
SoS Utility Health Check Options	229
Sample Output	230
VMware Cloud Builder Log Files	227
Getting Started with SDDC Manager	235
Log in to the SDDC Manager User Interface	236
Guided SDDC Manager Onboarding	236
Tour of the SDDC Manager User Interface	237
Navigation Bar	237
Log out of the SDDC Manager User Interface	239
Configure the Customer Experience Improvement Program Settings for VMware Cloud Foundation	239
Managing Certificates in VMware Cloud Foundation	240
View Certificate Information	240
Configure VMware Cloud Foundation to Use Microsoft CA-Signed Certificates	241
Prepare Your Microsoft Certificate Authority to Allow SDDC Manager to Manage Certificates	241
Configure a Microsoft Certificate Authority in SDDC Manager	245
Install Microsoft CA-Signed Certificates using SDDC Manager	246
Configure VMware Cloud Foundation to Use OpenSSL CA-Signed Certificates	247
Configure OpenSSL-signed Certificates in SDDC Manager	247
Install OpenSSL-signed Certificates using SDDC Manager	248
Install Third-Party CA-Signed Certificates Using Server Certificate and Certificate Authority Files	250
Install Third-Party CA-Signed Certificates in VMware Cloud Foundation Using a Certificate Bundle	252
Remove Old or Unused Certificates from SDDC Manager	255
Managing License Keys in VMware Cloud Foundation	255
Add a Component License Key in the SDDC Manager UI	256

Edit a Component License Key Description in the SDDC Manager UI	256
Delete a Component License Key in the SDDC Manager UI	257
Update Component License Keys for Workload Domain Components	257
ESXi Lockdown Mode	275
Managing Storage in VMware Cloud Foundation	282
Principal Storage	282
Supplemental Storage	283
vSAN Storage with VMware Cloud Foundation	283
Prerequisites for vSAN Storage	283
Procedures for vSAN Storage	284
vSAN Original Storage Architecture (OSA)	284
vSAN Express Storage Architecture (ESA)	284
vSAN Compute Clusters	284
Fibre Channel Storage with VMware Cloud Foundation	286
Prerequisites for FC Storage	286
Procedures for FC Storage	286
Sharing Remote Datastores with HCI Mesh for VI Workload Domains	655
Managing Workload Domains in VMware Cloud Foundation	311
About VI Workload Domains	312
Prerequisites for a Workload Domain	312
Creating VxRail VI Workload Domains	658
Create a VxRail VI Workload Domain in the SDDC Manager UI	659
Create a VxRail VI Workload Domain Using the Workflow Optimization Script	665
Delete a VI Workload Domain	317
View Workload Domain Details	317
Expand a Workload Domain	668
Add a VxRail Cluster to a Workload Domain Using the SDDC Manager UI	668
Add VxRail Hosts to a Cluster in VMware Cloud Foundation	671
Reduce a Workload Domain	672
Remove a Host from a Cluster in a Workload Domain	672
Delete a VxRail Cluster	672
Rename a Workload Domain	327
vSphere Cluster Management	328
View vSphere Cluster Details	328
Rename a Cluster in the SDDC Manager UI	328
Tag Management	330
Tag a Workload Domain	330
Tag a Cluster	331
Tag a Host	332
Managing NSX Edge Clusters in VMware Cloud Foundation	334

Prerequisites for an NSX Edge Cluster	335
Deploy an NSX Edge Cluster	335
Add Edge Nodes to an NSX Edge Cluster	341
Remove Edge Nodes from an NSX Edge Cluster	345
Managing Avi Load Balancer in VMware Cloud Foundation	346
Limitations of Avi Load Balancer in VMware Cloud Foundation	347
Deploy Avi Load Balancer for a Workload Domain	347
Remove Avi Load Balancer from a Workload Domain	350
Deploying Application Virtual Networks in VMware Cloud Foundation	351
Overlay-Backed NSX Segments	351
VLAN-Backed NSX Segments	351
VMware Aria Suite Components and NSX Segments	351
Identity and Access Management for VMware Cloud Foundation	352
Deploy Overlay-Backed NSX Segments	352
Deploy VLAN-Backed NSX Segments	353
VMware Cloud Foundation with VMware Tanzu	354
Enable Workload Management	355
View Workload Management Cluster Details	356
Update Workload Management License	356
VMware Aria Suite Lifecycle in VMware Cloud Foundation mode	356
VMware Aria Suite Lifecycle Implementation	357
Deploy VMware Aria Suite Lifecycle	358
Replace the Certificate of the VMware Aria Suite Lifecycle Instance	359
Configure Data Center and vCenter Server in VMware Aria Suite Lifecycle	359
Workspace ONE Access Implementation	360
Import the Workspace ONE Access Certificate to VMware Aria Suite Lifecycle	361
Add Workspace ONE Access Passwords to VMware Aria Suite Lifecycle	362
Deploy a Standard Workspace ONE Access Instance Using VMware Aria Suite Lifecycle	363
Deploy Clustered Workspace ONE Access Instance Using VMware Aria Suite Lifecycle	365
Configure an Anti-Affinity Rule and a Virtual Machine Group for a Clustered Workspace ONE Access Instance	367
Configure NTP on Workspace ONE Access	368
Configure the Domain and Domain Search Parameters on Workspace ONE Access	368
Configure an Identity Source for Workspace ONE Access	369
Add the Clustered Workspace ONE Access Cluster Nodes as Identity Provider Connectors	370
Assign Roles to Active Directory Groups for Workspace ONE Access	371
Assign Roles to Active Directory Groups for VMware Aria Suite Lifecycle	371
Working with NSX Federation in VMware Cloud Foundation	371
NSX Federation Key Concepts	372
NSX Federation Systems: Global Manager and Local Manager	372

NSX Federation Span: Local and Cross-Instance	372
NSX Federation Tunnel Endpoints.....	372
NSX Federation Tier Gateways.....	372
Configuring NSX Federation in VMware Cloud Foundation	373
Create Global Manager Clusters for VMware Cloud Foundation.....	374
Prepare Local Manager for NSX Federation in VMware Cloud Foundation	377
Enable NSX Federation in VMware Cloud Foundation.....	377
Stretch Segments between VMware Cloud Foundation Instances.....	379
Set Standby Global Manager	382
Replacing Global Manager Cluster Certificates in VMware Cloud Foundation	382
Import a CA-Signed Certificate to the Global Manager Cluster	383
Replace the Certificate for the First Global Manager Node	383
Replace Certificates and Virtual IP for the Remaining Global Manager Nodes.....	384
Update Local Manager Certificate Thumbprint in Global Manager Cluster	386
Password Management for NSX Global Manager Cluster in VMware Cloud Foundation.....	386
Backup and Restore of NSX Global Manager Cluster in VMware Cloud Foundation.....	386
Configure NSX Global Manager Cluster Backups	387
Restore an NSX Global Manager Cluster Backup	387
Stretching vSAN Clusters in VMware Cloud Foundation on Dell VxRail	730
About Availability Zones and Regions	731
Availability Zones	731
Regions	731
Stretched Cluster Requirements	395
VLANs and Subnets for Multiple Available Zones.....	395
Networking for Multiple Availability Zones.....	396
Deploy and Configure vSAN Witness Host	397
Deploy vSAN Witness Host.....	397
Configure the Management Network on the vSAN Witness Host.....	734
Register vSAN Witness Host	398
Configure NTP on the Witness Host	399
Configure the VMkernel Adapters on the vSAN Witness Host.....	399
Stretch a VxRail Cluster in VMware Cloud Foundation.....	736
NSX Configuration for Availability Zone 2.....	416
Configure IP Prefixes in the Tier-0 Gateway for Availability Zone 2.....	416
Configure Route Maps in the Tier-0 Gateway for Availability Zone 2.....	417
Configure BGP in the Tier-0 Gateway for Availability Zone 2.....	418
Configure Witness Traffic Separation for VMware Cloud Foundation on Dell VxRail	743
Create Distributed Port Groups for Witness Traffic.....	743
Delete Routes to the Witness Host	743
Add VMkernel Adapters for Witness Traffic.....	744

Configure the VMkernel Adapters for Witness Traffic	745
Expand a Stretched VxRail Cluster	745
Replace a Failed Host in a Stretched VxRail Cluster	746
Monitoring Capabilities in the VMware Cloud Foundation System	425
Viewing Tasks and Task Details.....	425
Viewing and Filtering Task Details	425
Managing Tasks and Subtask Details	426
Resizing the Task Panel.....	426
API Activity Logging.....	426
Activity Log Structure	426
Activity Log Example.....	426
Activity Logs Retention Policy	427
Log Analysis.....	427
Updating VMware Cloud Foundation DNS and NTP Servers.....	427
Update DNS Server Configuration	427
Update NTP Server Configuration.....	428
Supportability and Serviceability (SoS) Utility	429
SoS Utility Options.....	429
SoS Utility Help Options	429
SoS Utility Generic Options	430
SoS Utility VMware Cloud Foundation Summary Options	431
SoS Utility Fix-It-Up Options	431
SoS Utility Health Check Options.....	433
Example Health Check Commands:	434
Collect Logs for Your VMware Cloud Foundation System.....	434
Component Log Files Collected by the SoS Utility	437
Managing Users and Groups in VMware Cloud Foundation	441
Configuring the Identity Provider for VMware Cloud Foundation	442
Add Active Directory over LDAP or OpenLDAP as an Identity Source for VMware Cloud Foundation	442
Configure Microsoft ADFS as the Identity Provider in the SDDC Manager UI.....	444
Add a User or Group to VMware Cloud Foundation.....	446
Remove a User or Group	447
Create a Local Account	447
Create an Automation Account.....	449
Managing Passwords in VMware Cloud Foundation.....	772
Password Expiration Notifications	772
Rotate Passwords	454
Manually Update Passwords	456
Remediate Passwords	456
Look Up Account Credentials	457

Updating SDDC Manager Passwords	458
Update SDDC Manager Root and Super User Passwords.....	458
Update SDDC Manager Local Account Password	458
Update Expired SDDC Manager Root Password	459
Backing Up and Restoring SDDC Manager and NSX Manager	779
Reconfigure SFTP Backups for SDDC Manager and NSX Manager	460
File-Based Backups for SDDC Manager and vCenter Server.....	461
Back Up SDDC Manager	462
Configure a Backup Schedule for vCenter Server	463
Manually Back Up vCenter Server	463
Export the Configuration of the vSphere Distributed Switches	464
File-Based Restore for SDDC Manager, vCenter Server, and NSX.....	465
Restore SDDC Manager	465
Restore vCenter Server	468
Restore the Configuration of a vSphere Distributed Switch	474
Restore an NSX Manager Cluster Node	474
Restoring NSX Edge Cluster Nodes	483
Image-Based Backup and Restore of VMware Cloud Foundation	490
Upgrading to VMware Cloud Foundation 5.2.x on Dell VxRail	810
SDDC Manager Functionality During an Upgrade to VMware Cloud Foundation 5.2.....	499
Upgrade States and Terminology	499
vSphere UI Client Plug-ins	502
Monitor VMware Cloud Foundation Updates	502
View VMware Cloud Foundation Update History	504
Access VMware Cloud Foundation Upgrade Log Files.....	504
Downloading VMware Cloud Foundation Upgrade Bundles	505
Online and Offline Downloads.....	505
Other Bundle Types	505
Connect SDDC Manager to a Software Depot for Downloading Bundles	389
Download Bundles Using SDDC Manager	507
Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles	508
Offline Download of Independent SDDC Manager Bundles	511
Offline Download of Async Patch Bundles	513
Offline Download of Flexible BOM Upgrade Bundles	515
HCL Offline Download for VMware Cloud Foundation	517
Download Bundles to an Offline Depot	519
VMware Cloud Foundation Upgrade Prerequisites	520
VMware Cloud Foundation 5.2.x Upgrade Overview	521
VMware Cloud Foundation Upgrade Preparation	521
Management Domain Upgrade	521

VI Workload Domain Upgrade	524
Upgrade the Management Domain to VMware Cloud Foundation 5.2.x.....	526
Perform Update Precheck - Versions Prior to SDDC Manager 5.0.....	527
Perform Update Precheck in SDDC Manager	529
Apply the VMware Cloud Foundation 5.2.x Upgrade Bundle.....	533
Apply VMware Cloud Foundation Configuration Updates.....	534
Upgrade VMware Aria Suite Lifecycle and VMware Aria Suite Products for VMware Cloud Foundation.....	538
Upgrade NSX for VMware Cloud Foundation in a Federated Environment.....	538
Upgrade NSX for VMware Cloud Foundation 5.2.x	540
Upgrade vCenter Server for VMware Cloud Foundation 5.2.x	542
Upgrade VxRail Manager and ESXi Hosts for VMware Cloud Foundation.....	857
Upgrade vSAN Witness Host for VMware Cloud Foundation	547
Upgrade vSphere Distributed Switch versions.....	559
Upgrade vSAN on-disk format versions	559
Update License Keys for a Workload Domain	558
Upgrade VI Workload Domains to VMware Cloud Foundation 5.2.x	560
Plan VI Workload Domain Upgrade	560
Perform Update Precheck in SDDC Manager	529
Upgrade NSX for VMware Cloud Foundation in a Federated Environment.....	538
Upgrade NSX for VMware Cloud Foundation 5.2.x	540
Upgrade vCenter Server for VMware Cloud Foundation 5.2.x	542
Upgrade VxRail Manager and ESXi Hosts for VMware Cloud Foundation.....	857
Upgrade vSAN Witness Host for VMware Cloud Foundation	547
Upgrade vSphere Distributed Switch versions.....	559
Upgrade vSAN on-disk format versions	559
Update License Keys for a Workload Domain	558
Independent SDDC Manager Upgrade using the SDDC Manager UI.....	584
Flexible BOM Upgrade in VMware Cloud Foundation.....	584
Patching the Management and Workload Domains	587
Troubleshooting for Upgrading VMware Cloud Foundation	588
SDDC Manager Troubleshooting	588
vCenter Troubleshooting	589
Shutdown and Startup of VMware Cloud Foundation.....	494
Shutting Down VMware Cloud Foundation	495
Starting Up VMware Cloud Foundation	495
Instance Recovery Guide	878
Example Failure Scenarios	878
Intended Audience	878
Related VMware Cloud Foundation Documentation	878
Verifications and Remediations.....	878

Check SDDC Manager Health.....	878
Delete the Temporary vCenter Server Instance	878
Perform SDDC Manager Pre-Checks.....	879
Introducing Security and Compliance for VMware Cloud Foundation 5.2	880
Intended Audience	880
Required VMware Software.....	880
Security by Design.....	880
Security Architecture.....	881
Security Principles.....	881
Governance, Risk, and Compliance and Mapping.....	882
Control Definition	883
Cybersecurity Considerations.....	883
Business Impact Assessment.....	883
Compliance Kit for VMware Cloud Foundation.....	883
<i>Compliance Kit for VMware Cloud Foundation Structure</i>	<i>883</i>
VMware Cloud Foundation Compliance Kit.....	884
Default Access Controls Configured in VMware Cloud Foundation.....	884
Security and Compliance Configuration Guide for VMware Cloud Foundation 5.2	888
Intended Audience	888
Required VMware Software.....	888
Update History	888
Software Requirements.....	888
Securing ESXi Hosts	890
Security Best Practices for Securing ESXi Hosts.....	890
Configure Multiple Security Settings on the ESXi Hosts by Using the ESXi Shell	890
Configure Multiple Security Settings on the ESXi Hosts by Using PowerCLI	891
Configure Multiple Security Settings on Unassigned ESXi Hosts by Using PowerCLI	896
Activate Normal Lockdown Mode on the ESXi Hosts.....	896
Securing vCenter Server	897
Security Best Practices for Securing vCenter Server	897
Configure Security Settings for vCenter Server from the vSphere Client.....	899
Configure Security Settings for vCenter Server by Using PowerCLI.....	902
Configure Security Settings on the vCenter Server Appliance	904
Securing SDDC Manager.....	904
Security Best Practices for Securing SDDC Manager.....	904
Configure Security Settings for SDDC Manager by Using the SDDC Manager UI	906
Securing Management Virtual Machines	907
Securing vSAN	908
Security Best Practices for Securing vSAN.....	908
Configure a Proxy Server for vSAN from the vSphere Client.....	909

Securing VMware NSX	909
Security Best Practices for Securing VMware NSX	909
Configure Security Settings for VMware NSX by Using the User Interfaces.....	910
Configure Security Settings for NSX by Using CLI Commands	911
Configure Security Settings for VMware NSX by Using NSX API	912
Optional Security Configurations for VMware NSX	912
Configure Security Settings for NSX Edge Nodes by Using the User Interface	912
Configure Security Settings for NSX Edge Nodes by Using CLI Commands	915
Configure Security Settings for Distributed Firewall by Using the User Interface	916
Security Configurations Not Applicable or Not Compatible with VMware Cloud Foundation	917
Documentation Legal Notice	919

Release Notes

The *Release Notes* for VMware Cloud Foundation and VMware Cloud Foundation on Dell EMC VxRail provide information about each release, including: What's new in the release, software components and versions included in the Bill of Materials (BOM), resolved issues, and known issues.

VMware Cloud Foundation 5.2.1 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Available Languages](#)
- [Important VMware Cloud Foundation 5.2.1 API Changes](#)
- [Deprecation Notices](#)
- [VMware Cloud Foundation Bill of Materials \(BOM\)](#)
- [Supported Hardware](#)
- [Documentation](#)
- [Browser Compatibility and Screen Resolutions](#)
- [Installation and Upgrade Information](#)
- [VMware Cloud Foundation 5.2.1.1 Release Information](#)
- [Resolved Issues](#)
- [Known Issues](#)

Introduction

VMware Cloud Foundation 5.2.1 | 09 OCT 2024 | Build 24392123

VMware Cloud Foundation 5.2.1.1 | 05 DEC 2024 | Build 24397777

Check for additions and updates to these release notes.

What's New

The VMware Cloud Foundation (VCF) 5.2.1 release includes the following:

- **Reduced Downtime Upgrade (RDU) support for vCenter:** VCF users can now leverage vCenter Reduced Downtime Upgrade (RDU) to execute a vCenter upgrade. vCenter RDU is a migration-based approach to upgrading vCenter and reduces the vCenter downtime to less than 5 minutes.
- **NSX in-place upgrades for clusters that use vSphere Lifecycle Manager baselines:** VCF users now have the choice to perform NSX in-place upgrade for clusters that use vSphere Lifecycle Manager baselines. In-place upgrades eliminate the need to place hosts into maintenance mode during the upgrade.
- **Support for vSphere Lifecycle Manager baseline and vSphere Lifecycle Manager image-based clusters in same workload domain:** VCF users now have the flexibility to deploy and upgrade vLCM baseline and vLCM image-based clusters within the same workload domain. With this capability, VCF users can now deploy vSAN ESA clusters and vSAN OSA clusters in the same workload domain.
- **Support for the "License Now" option for vSAN add-on licenses based on capacity per tebibyte (TiB):** VCF users can now apply the vSAN TiB capacity license within the SDDC Manager UI to expand storage capacity for their workload domains and clusters. You can also use the "License Later" option to assign the per-TiB vSAN license key using the vSphere Client.
- **Set up VMware Private AI Foundation infrastructure from the vSphere Client:** VCF users can leverage a new guided workflow in the vSphere Client to set up infrastructure for VMware Private AI Foundation and maximize the

potential of NVIDIA GPU-enabled ESXi hosts. The workflow streamlines the set up process by centralizing configuration steps from SDDC Manager and vCenter into a single workflow.

- **Manage all SDDC certificates and passwords from a single UI:** SDDC Manager certificate and password management functionality is now integrated in the vSphere Client to simplify and speed-up day-to-day operations. VCF users can now manage the certificates, integrated certificate authorities, and system user passwords from the Administration section in the vSphere Client.

Available Languages

Beginning with the next major release, VCF will be supporting the following localization languages:

- Japanese
- Spanish
- French

The following languages will no longer be supported:

- Italian, German, and Simplified Chinese.

Impact:

- Customers who have been using the deprecated languages will no longer receive updates in these languages.
- All user interfaces and help documentation will be available only in English or in the three supported languages mentioned above.

Because VCF localization utilizes the browser language settings, ensure that your settings match the desired language.

Important VMware Cloud Foundation 5.2.1 API Changes

Change	Impacted APIs
Fixed usage of wildcard ("*/") if media type is not defined for API response metadata.	<ul style="list-style-type: none"> • POST /v1/sddc-manager/trusted-certificates • POST /v1/resources/license-check • POST /v1/identity-providers • POST /v1/identity-providers/{id}/identity-sources • PATCH /v1/system/proxy-configuration • PATCH /v1/sddcs/validations/{id} • DELETE /v1/identity-providers/{id} • DELETE /v1/identity-providers/{id}/identity-sources/{domainName} • PATCH /v1/identity-providers/{id}/identity-sources/{domainName} • GET /v1/resources/license-checks/{id} • GET /v1/licensing-info • GET /v1/licensing-info/system • GET /v1/licensing-info/domains/{id}
Fixed ambiguous/empty response content type in API response metadata.	<ul style="list-style-type: none"> • GET /v1/clusters/{id}/network/criteria • GET /v1/sddc-manager/upgradables • PUT /v1/compatibility-matrices
Fixed responses having schema type as a string instead of having respective user defined type.	<ul style="list-style-type: none"> • POST /v1/hosts • GET /v1/hosts
Response type for 4xx/5xx responses has been fixed to ErrorResponse type.	<ul style="list-style-type: none"> • GET /v1/system/ntp-configuration/validations • GET /v1/system/dns-configuration/validations

Table continued on next page

Continued from previous page

Change	Impacted APIs
	<ul style="list-style-type: none"> • POST /v1/vasa-providers/validations • POST /v1/hosts/validations • POST /v1/hosts/validations/commissions • GET /v1/sddcs/validations/{id} • GET /v1/vasa-providers/validations/{id} • GET /v1/hosts/validations/{id} • GET /v1/edge-clusters/validations/{id} • DELETE /v1/host • PATCH /v1/system/proxy-configuration • GET /v1/system/proxy-configuration • GET /v1/system/security/fips • GET /v1/sddcs/{id}/detail-report • GET /v1/sddcs/validations/{validationId}/report
Fixed media-type for error responses.	<ul style="list-style-type: none"> • GET /v1/sddcs/{id}/detail-report • GET /v1/sddcs/validations/{validationId}/report
Removed incorrect and unused error API response definitions for 404 and 501 error codes from the OpenAPI specification.	<p>404 is removed from:</p> <ul style="list-style-type: none"> • GET /v1/sddcs/about • GET /v1/sddcs/validations <p>501 is removed from:</p> <ul style="list-style-type: none"> • GET /v1/sddcs/about • GET v1/sddcs/validations/{id} • GET /v1/sddcs/validations • GET /v1/sddcs/{id}/sddc-manager • POST /v1/system/sddc-spec-converter

Deprecation Notices

In a future major release, the following APIs will not be supported:

- POST /v1/bundles
- POST /v1/product-version-catalog

The APIs will be replaced with:

- POST /v1/product-binaries
- POST /v1/product-version-catalogs

After upgrading SDDC Manager to this new release, automation customers who use the unsupported APIs should transition to the new APIs.

- The following features are being deprecated and will be removed in a future major release:

- Cloud Builder Appliance
- Cloud Builder APIs
- Cloud Builder deployment parameter workbooks
- NSX Edge management workflow

- VMware End Of Availability of Perpetual Licensing and SaaS Services. See <https://blogs.vmware.com/cloud-foundation/2024/01/22/vmware-end-of-availability-of-perpetual-licensing-and-saas-services/> for more information.

- In a future release, the "Connect Workload Domains" option from the VMware Aria Operations card located in **SDDC Manager > Administration > Aria Suite** section will be removed and related VCF Public API options will be deprecated.

Starting with VMware Aria Operations 8.10, functionality for connecting VCF Workload Domains to VMware Aria Operations is available directly from the UI. Users are encouraged to use this method within the VMware Aria Operations UI for connecting VCF workload domains, even if the integration was originally set up using SDDC Manager.

- Deprecation announcements for VMware NSX. See the [VMware NSX 4.2.1 Release Notes](#) for details.

VMware Cloud Foundation Bill of Materials (BOM)

The VMware Cloud Foundation software product is comprised of the following software Bill-of-Materials (BOM). The components in the BOM are interoperable and compatible.

Software Component	Version	Date	Build Number
Cloud Builder VM	5.2.1	09 OCT 2024	24307856
SDDC Manager	5.2.1	09 OCT 2024	24307856
VMware vCenter Server Appliance	8.0 Update 3c	09 OCT 2024	24305161
VMware ESXi	8.0 Update 3b	17 SEP 2024	24280767
VMware vSAN Witness Appliance	8.0 Update 3	19 JUN 2024	24022510
VMware NSX	4.2.1	09 OCT 2024	24304122
VMware Aria Suite Lifecycle	8.18	23 JUL 2024	24029603

- VMware vSAN is included in the VMware ESXi bundle.
- You can use VMware Aria Suite Lifecycle to deploy VMware Aria Automation, VMware Aria Operations, VMware Aria Operations for Logs, and Workspace ONE Access. VMware Aria Suite Lifecycle determines which versions of these products are compatible and only allows you to install/upgrade to supported versions.
- VMware Aria Operations for Logs content packs are installed when you deploy VMware Aria Operations for Logs.
- The VMware Aria Operations management pack is installed when you deploy VMware Aria Operations.
- You can access the latest versions of the content packs for VMware Aria Operations for Logs from the VMware Solution Exchange and the VMware Aria Operations for Logs in-product marketplace store.

Supported Hardware

For details on supported configurations, see the [VMware Compatibility Guide \(VCG\)](#) and the Hardware Requirements section on the Prerequisite Checklist tab in the [Planning and Preparation Workbook](#).

Documentation

To access the VCF documentation, go to the [VMware Cloud Foundation product documentation](#).

To access the documentation for VMware software products that SDDC Manager can deploy, see the product documentation and use the drop-down menus on the page to choose the appropriate version:

- [VMware vSphere product documentation](#), includes the documentation for ESXi and vCenter Server
- [VMware vSAN product documentation](#)
- [VMware NSX product documentation](#)

Browser Compatibility and Screen Resolutions

The VMware Cloud Foundation web-based interface supports the latest two versions of the following web browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge

For the Web-based user interfaces, the supported standard resolution is 1920 by 1080 pixels.

Installation and Upgrade Information

You can install VMware Cloud Foundation 5.2.1 as a new release or perform a sequential or skip-level upgrade to VMware Cloud Foundation 5.2.1.

Installing as a New Release

The new installation process has three phases:

- **Phase One: Prepare the Environment:** The [Planning and Preparation Workbook](#) provides detailed information about the software, tools, and external services that are required to implement a Software-Defined Data Center (SDDC) with VMware Cloud Foundation, using a standard architecture model.
- **Phase Two: Image all servers with ESXi:** Image all servers with the ESXi version mentioned in the Cloud Foundation Bill of Materials (BOM) section. See the [VMware Cloud Foundation Deployment Guide](#) for information on installing ESXi.
- **Phase Three: Install Cloud Foundation 5.2.1:** See the [VMware Cloud Foundation Deployment Guide](#) for information on deploying Cloud Foundation.

Upgrading to Cloud Foundation 5.2.1

You can perform a sequential or skip-level upgrade to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 4.5.0 or later. If your environment is at a version earlier than 4.5.0, you must upgrade the management domain and all VI workload domains to VMware Cloud Foundation 4.5.0 or above and then upgrade to VMware Cloud Foundation 5.2.1. For more information see [VMware Cloud Foundation Lifecycle Management](#).

IMPORTANT

Before you upgrade a vCenter Server, take a file-based backup. See [Manually Back Up vCenter Server](#).

NOTE

Since VMware Cloud Foundation disables the SSH service by default, scripts that rely on SSH being enabled on ESXi hosts will not work after upgrading to VMware Cloud Foundation 5.2.1. Update your scripts to account for this new behavior. See [KB 86230](#) for information about enabling and disabling the SSH service on ESXi hosts.

VMware Cloud Foundation 5.2.1.1 Release Information

VMware Cloud Foundation 5.2.1.1 includes bug fixes and a new version of SDDC Manager. In addition, using the VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1 adds support for a broader range of vSphere environments and topologies and relaxes some of the constraints that were in place with the previous versions. Improvements include:

- The ability to import vSphere clusters that have shared vSphere Distributed Switches (VDS)
- Support for importing clusters with LACP enabled
- Importing vSphere Environments with a mix of vLCM Images and Baselines
- Relaxing of Guardrails for vSphere Standard Switches and Standalone Hosts

NOTE: The VCF Import Tool 5.2.1.1 has been replaced with the VCF Import Tool 5.2.1.2. Do not use the VCF Import Tool 5.2.1.1.

You can upgrade to VMware Cloud Foundation 5.2.1.1 from VMware Cloud Foundation 5.2 or later.

Upgrading from 5.2.1:

SDDC Manager is the only component that requires an upgrade. See [Independent SDDC Manager Upgrade using the SDDC Manager UI](#).

Upgrading from 5.2:

See [Flexible BOM Upgrade in VMware Cloud Foundation](#). When selecting the target version for SDDC Manager choose the version listed in the BOM update table below.

Known issues:

- In order to upgrade from 5.2 to 5.2.1.1, you must download the bundles for both SDDC Manager 5.2.1.0 and SDDC Manager 5.2.1.1.
- The Bundle Management window in the SDDC Manager UI displays "VMware Cloud Foundation Update 5.2.1.0" instead of "VMware Cloud Foundation Update 5.2.1.1" for the 5.2.1.1 bundle. The description of the bundle correctly describes it as the upgrade bundle for 5.2.1.1. This is a cosmetic issue only and does not impact the upgrade.

VMware Cloud Foundation 5.2.1.1 contains the following BOM updates:

Software Component	Version	Date	Build Number
SDDC Manager	5.2.1.1	05 DEC 2024	24397777

Resolved Issues

The following issues are resolved in this release:

- VMware Cloud Foundation 5.2 does not support the "License Now" option for vSAN add-on licenses based on capacity per terabyte (TiB).
- Remove unresponsive ESXi Host fails when SDDC Manager certificate does not have subject alternative name.

Known Issues

VMware Cloud Foundation Known Issues

VCF Import Tool does not support clusters that use vSphere Configuration Profiles

If you use the VCF Import Tool to import/convert an existing vSphere environment that includes clusters that use vSphere Configuration Profiles, the task fails during NSX deployment.

Workaround: None. Clusters that use vSphere Configuration Profiles do not support NSX.

Primary datastore is not getting set for imported workload domains with NFS 4.1 datastore

When you use the VCF Import Tool to import a cluster for which NFS 4.1 is the only shared datastore, the primary datastore and datastore type is not getting set in VCF and the workload domain is not visible in the SDDC Manager UI. See <https://knowledge.broadcom.com/external/article/372424> for details.

Workaround: None.

Limitations for importing vSAN clusters

When you use the VCF Import Tool to import a vSAN cluster, you should avoid importing clusters with certain configurations. SDDC Manager day-N operations will not be supported on imported vSAN clusters with these configurations. See <https://knowledge.broadcom.com/external/article/371494> for details.

Workaround: None.

Lifecycle Management Precheck does not throw an error when NSX Manager inventory is out of sync

The Lifecycle Management Precheck displays a green status and does not generate any errors for NSX Manager inventory.

Workaround: None

Upgrade Pre-Check Scope dropdown may contain additional entries

When performing Upgrade Prechecks through SDDC Manager UI and selecting a target VCF version, the Pre-Check Scope dropdown may contain more selectable entries than necessary. SDDC Manager may appear as an entry more than once. It also may be included as a selectable component for VI domains, although it's a component of the management domain.

Workaround: None. The issue is visual with no functional impact.

Converting clusters from vSphere Lifecycle Manager baselines to vSphere Lifecycle Manager images is not supported.

vSphere Lifecycle Manager baselines (previously known as vSphere Update Manager or VUM) are deprecated in vSphere 8.0, but continue to be supported. See [KB article 89519](#) for more information.

VMware Cloud Foundation does not support converting clusters from vSphere Lifecycle Manager baselines to vSphere Lifecycle Manager images. This capability will be supported in a future release.

Workaround: None

Upgrade Known Issues

Upgrade precheck warning "ESXi upgrade policy validation across vCenter and SDDC Manager"

In SDDC Manager, the default upgrade policy applies to all clusters, while in vSphere each cluster has a distinct upgrade policy. This can create a scenario where the ESX upgrade policy configured in vCenter does not match what SDDC Manager expects.

Workaround: This issue causes a warning only. You can proceed with the upgrade without remediating the issue.

Incorrect backup options displayed for a vCenter Regular Update

When you configure a vCenter upgrade in the SDDC Manager UI, the Backup screen shows the options for a Reduced Downtime Upgrade (RDU), even if you selected vCenter Regular Update as the update mechanism. Do not select "I wish to continue without a backup of the vCenter server" when performing a vCenter Regular Update.

Workaround: None. For a vCenter Regular Update, you must back up vCenter before you upgrade.

Bundle Transfer Utility fails to upload the NSX Advanced Load Balancer install bundle

If you on a pre-5.2.x version of VMware Cloud Foundation and use the Bundle Transfer Utility to download all bundles for VCF 5.2.x, then uploading the NSX Advanced Load Balancer install bundle fails. This bundle is only supported with SDDC Manager 5.2 and later.

Workaround: Upgrade SDDC Manager to 5.2 or later and then retry uploading the NSX Advanced Load Balancer install bundle.

NSX host cluster upgrade fails

If you are upgrading a workload domain that uses vSphere Lifecycle Manager images and its cluster image was created from an ESXi host that uses vSphere Lifecycle Manager baselines, then NSX host cluster upgrade will fail. A cluster image created from an ESXi host that uses vSphere Lifecycle Manager baselines contains an NSX component that causes this issue.

NOTE: This issue is resolved in you have ESXi and vCenter Server 8.0 Update 3 or later.

Workaround: Do not create cluster images from an ESXi host that uses vSphere Lifecycle Manager baselines. If you encounter this issue, you can resolve it by using the vSphere Client to remove the NSX LCP Bundle component from the

cluster image.

Summary Monitor Configure Permissions Hosts VMs Datastores Networks Updates

Hosts **Edit Image**
Select the version of ESXi and other components that you want for the hosts in this cluster. The same image will be applied consistently to all these hosts.

Image
Hardware Compatibility
VMware Tools
VM Hardware

ESXi Version [8.0 Update 2 - 23718068](#) (released 04/21/2024)

Vendor Addon [SELECT](#) (optional)

Firmware and Drivers Addon [SELECT](#) (optional)

Components [1 additional components](#) [Hide details](#)

[ADD COMPONENTS](#) Show [Additional components](#)

Component Name	Version	Notes
NSX LCP Bundle	NSX LCP Bundle(4.1.2.0.0-8.0.22305537)	Manually added component

SDDC Manager UI shows the incorrect source version when upgrading SDDC Manager

When you view the VMware Cloud Foundation Update Status for SDDC Manager, the UI may show the incorrect source version.

Updating [VIEW UPDATE ACTIVITY](#)

▼ SDDC MANAGER - VMware Cloud Foundation Update 5.2.0.0 5.11.0-23480823 → 5.2.0.0-24026533

Workaround: None. This is a cosmetic issue only and does not affect the upgrade.

Workspace ONE Access inventory sync fails in SDDC Manager after upgrading VMware Aria Suite Lifecycle

After upgrading Aria Suite Lifecycle to version 8.12 or later, triggering a Workspace ONE Access inventory sync from Aria Suite Lifecycle fails. The SDDC Manager UI reports the following error: Failed to configure WSA <wsa_fqdn> in vROps .vrops_fqdn>, because Failed to manage vROps adapter.

Workaround: Download the bundle for your version of Aria Suite Lifecycle to SDDC Manager and retry the inventory sync.

VCF ESXi upgrade fails during post validation due to HA related cluster configuration issue

The upgrade of ESXi Cluster fails with error that is similar to below error message:

```
Cluster Configuration Issue: vSphere HA failover operation in progress in cluster
<cluster-name> in datacenter <datacenter-name>: 0 VMs being restarted, 1 VMs waiting for a
retry, 0 VMs waiting for resources, 0 inaccessible vSAN VMs
```

Workaround: See [KB article 90985](#).

Lifecycle Management Precheck does not throw an error when NSX Manager inventory is out of sync

Workaround None.

NSX upgrade may fail if there are any active alarms in NSX Manager

If there are any active alarms in NSX Manager, the NSX upgrade may fail.

Workaround: Check the NSX Manager UI for active alarms prior to NSX upgrade and resolve them, if any. If the alarms are not resolved, the NSX upgrade will fail. The upgrade can be retried once the alarms are resolved.

SDDC Manager upgrade fails at "Setup Common Appliance Platform"

If a virtual machine reconfiguration task (for example, removing a snapshot or running a backup) is taking place in the management domain at the same time you are upgrading SDDC Manager, the upgrade may fail.

Workaround: Schedule SDDC Manager upgrades for a time when no virtual machine reconfiguration tasks are happening in the management domain. If you encounter this issue, wait for the other tasks to complete and then retry the upgrade.

Parallel upgrades of vCenter Server are not supported

If you attempt to upgrade vCenter Server for multiple VI workload domains at the same time, the upgrade may fail while changing the permissions for the vpostgres configuration directory in the appliance. The message `chown -R vpostgres:vpgmongrp /storage/archive/vpostgres` appears in the PatchRunner.log file on the vCenter Server Appliance.

Workaround: Each vCenter Server instance must be upgraded separately.

When you upgrade VMware Cloud Foundation, one of the vSphere Cluster Services (vCLS) agent VMs gets placed on local storage

vSphere Cluster Services (vCLS) ensures that cluster services remain available, even when the vCenter Server is unavailable. vCLS deploys three vCLS agent virtual machines to maintain cluster services health. When you upgrade VMware Cloud Foundation, one of the vCLS VMs may get placed on local storage instead of shared storage. This could cause issues if you delete the ESXi host on which the VM is stored.

Workaround: Deactivate and reactivate vCLS on the cluster to deploy all the vCLS agent VMs to shared storage.

1. Check the placement of the vCLS agent VMs for each cluster in your environment.
 - a. In the vSphere Client, select **Menu > VMs and Templates**.
 - b. Expand the vCLS folder.
 - c. Select the first vCLS agent VM and click the Summary tab.
 - d. In the Related Objects section, check the datastore listed for Storage. It should be the vSAN datastore. If a vCLS agent VM is on local storage, you need to deactivate vCLS for the cluster and then re-enable it.
 - e. Repeat these steps for all vCLS agent VMs.

2. Deactivate vCLS for clusters that have vCLS agent VMs on local storage.
 - a. In the vSphere Client, click **Menu > Hosts and Clusters**.
 - b. Select a cluster that has a vCLS agent VM on local storage.
 - c. In the web browser address bar, note the moref id for the cluster.

For example, if the URL displays as `https://vcenter-1.vrack.vsphere.local/ui/app/cluster;nav=h/urn:vmomi:ClusterComputeResource:domain-c8:503a0d38-442a-446f-b283-d3611bf035fb/summary`, then the moref id is domain-c8.
 - d. Select the vCenter Server containing the cluster.
 - e. Click **Configure > Advanced Settings**.
 - f. Click **Edit Settings**.
 - g. Change the value for `config.vcls.clusters.<moref id>.enabled` to `false` and click **Save**.

If the `config.vcls.clusters.<moref id>.enabled` setting does not appear for your moref id, then enter its Name and `false` for the Value and click **Add**.
 - h. Wait a couple of minutes for the vCLS agent VMs to be powered off and deleted. You can monitor progress in the Recent Tasks pane.

3. Enable vCLS for the cluster to place the vCLS agent VMs on shared storage.
 - a. Select the vCenter Server containing the cluster and click **Configure > Advanced Settings**.
 - b. Click **Edit Settings**.
 - c. Change the value for `config.vcls.clusters.<moref id>.enabled` to `true` and click **Save**.
 - d. Wait a couple of minutes for the vCLS agent VMs to be deployed and powered on. You can monitor progress in the Recent Tasks pane.

4. Check the placement of the vCLS agent VMs to make sure they are all on shared storage

You are unable to update NSX Data Center in the management domain or in a workload domain with vSAN principal storage because of an error during the NSX transport node precheck stage

In SDDC Manager, when you run the upgrade precheck before updating NSX Data Center, the NSX transport node validation results with the following error.

No coredump target has been configured. Host core dumps cannot be saved.:System logs on host sfo01-m01-esx04.sfo.rainpole.io are stored on non-persistent storage. Consult product documentation to configure a syslog server or a scratch partition.

Because the upgrade precheck results with an error, you cannot proceed with updating the NSX Data Center instance in the domain. VMware Validated Design supports vSAN as the principal storage in the management domain. However, vSAN datastores do not support scratch partitions. See VMware Knowledge Base article [2074026](#).

Disable the update precheck validation for the subsequent NSX Data Center update.

1. Log in to SDDC Manager as **vcf** using a Secure Shell (SSH) client.
2. Open the `application-prod.properties` file for editing: `vi /opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties`
3. Add the following property and save the file: `lcm.nsxt.suppress.prechecks=true`
4. Restart the life cycle management service: `systemctl restart lcm`
5. Log in to the SDDC Manager user interface and proceed with the update of NSX Data Center.

Bring-up Known Issues

SDDC Manager Known Issues

The SDDC Manager UI displays incorrect CPU and memory utilization information for hosts

When you view CPU and memory usage for hosts in the SDDC Manager UI (**Inventory > Hosts**), the information may not reflect actual utilization.

FQDN	Host IP	Network Pool	Configuration Status	Host State	Cluster	CPU Usage	Memory Usage
esxi-1.vrack.vsphere.local	10.0.0.100	bringup-networkpool	Active	Assigned (sddcid-1001)	SDDC-Cluster1	4%	51%

Workaround: Use the vSphere Client to view CPU and memory utilization information for hosts.

Install bundle for VMware Aria Suite Lifecycle 8.18 displays incomplete information

The SDDC Manager UI displays incomplete information for the VMware Aria Suite Lifecycle 8.18 install bundle.

VMware Software Install Bundle - vRealize Suite Lifecycle Manager

8.18.0-24029603

Released Jun 14, 2024 2 GB

This VMware Software Upgrade contains [Add details here]. For more information, please see - [Insert link to release notes / KB article here]

[View Details](#)

Workaround: None. This is a cosmetic issue and does not impact your ability to download or use the bundle.

When creating a network pool, the IP addresses you provide are not validated to ensure that they are not in use SDDC Manager *does* validate the included IPs for a new network pool against other network pools and against all other network types (vSAN, NFS, and so on) being added to the new network pool. However, it *does not* validate them against other components that are already deployed in the VMware Cloud Foundation instance (for example, ESXi hosts, NSX Managers, and so on). This can result in duplicate IP address errors or failed workflows.

Workaround: When creating a network pool, do not include any IP addresses that are already in use. If you already created a network pool that includes IP addresses that are used by other components, contact Broadcom Support to resolve the issue.

vSphere Lifecycle Manager images that utilize “removed components” are not supported

Starting with vSphere 8.0 Update 3, you can remove the Host Client and VMware Tools components from a base image, remove unnecessary drivers from vendor add-ons and components, and override existing drivers in a desired image. SDDC Manager does not support this functionality yet for imported or extracted cluster images.

Workaround: None.

Workload Domain Known Issues

Deploying Avi Load Balancer fails

When you deploy Avi Load Balancer (formerly known as NSX Advanced Load Balancer), the deployment may fail with the message `OVA upload to NSX failed`. This can happen if the certificates of the management domain NSX Manager nodes do not include their IP addresses in their Subject Alternative Names (SANs).

Workaround: Generate new CSRs for the management domain NSX Manager nodes, making sure to include IP addresses for the SANs. For example:

Subject Alternative Name ×

Enter the SAN for each resource you selected.

SAN(s)	nsx-mgmt-1.vrack.vsphere.local
Optional	172.16.11.132
	E.g. nsx-mgmt-1.vrack.vsphere.local, 1, etc.
SAN(s)	vip-nsx-mgmt.vrack.vsphere.local
Optional	172.16.11.131
	E.g. vip-nsx-mgmt.vrack.vsphere.local, 1, etc.

Generate the signed certificates using the CSRs and then install the signed certificates in the NSX Manager nodes. See [Managing Certificates in VMware Cloud Foundation](#) for more information.

Once the new certificates are installed, retry deploying Avi Load Balancer.

Switch configuration error when deploying a VI workload domain or adding a cluster to a workload domain with hosts that have two DPUs

If you are using ESXi hosts with two data processing units (DPU) to deploy a new VI workload domain or add a cluster to a workload domain, you may see the following error during switch configuration: *Error in validating Config Profiles*. This can be caused by the presence of a vusb0 network adapter on the hosts.

Workaround: Contact Broadcom Support to remove the vusb0 interface from the SDDC Manager inventory.

Deploying a VI workload domain or adding a cluster to a workload domain fails with hosts that have two DPUs

If you are using ESXi hosts with two data processing units (DPU) to deploy a new VI workload domain or add a cluster to a workload domain, the task fails when adding the ESXi hosts to the vSphere Distributed Switch (VDS) with the error *Cannot complete a vSphere Distributed Switch operation for one or more host members*.

The VDS created by SDDC Manager for dual DPU hosts has all 4 uplinks in Active mode and this does not work with an NSX uplink profile where one set of DPU uplinks is Active and a second set of DPU uplinks is Standby.

Workaround: Use the vSphere Client to manually update the DPU failover settings for the VDS and then retry the workflow from SDDC Manager.

1. In the vSphere Client, browse to the VDS in the vCenter that contains the hosts.
2. Click the **Configure** tab and select **DPU Failover Settings**.

The screenshot shows the vds-nsx configuration interface. The 'Configure' tab is selected. Under 'Settings', 'DPU Failover Settings' is expanded. The 'Active uplinks' section lists uplink4, uplink3, uplink2, and uplink1. The 'Standby uplinks' section is currently empty. The 'DPU Failover Settings' menu item is highlighted.

Section	Uplink
Active uplinks	uplink4
	uplink3
	uplink2
	uplink1
Standby uplinks	

3. Click **Edit** and move uplink3 and uplink4 from Active to Standby.
4. Click **OK**.

The screenshot shows the vds-nsx configuration interface after the changes. The 'Active uplinks' section now lists uplink2 and uplink1. The 'Standby uplinks' section now lists uplink4 and uplink3. The 'DPU Failover Settings' menu item is highlighted.

Section	Uplink
Active uplinks	uplink2
	uplink1
Standby uplinks	
	uplink4
	uplink3

5. In the SDDC Manager UI, retry the failed workflow.

NSX Edge cluster deployment fails at "Create VLAN Port Group" stage with message "Invalid parameter: port group already exists"

When you deploy an NSX Edge cluster for VI workload domain and you select the option "USE ESXI MANAGEMENT VMK'S VLAN", the management portgroup name and VLAN ID are auto-populated. SDDC Manager tries to create a portgroup with same VLAN and portgroup name as ESXi management, but since the portgroup name already exists in vCenter the operation fails.

Workaround: If you select the option "USE ESXI MANAGEMENT VMK'S VLAN", change the auto-populated portgroup name to something else so that there is no conflict. If the environment is already in failed state, remove the partially deployed edge cluster. See <https://knowledge.broadcom.com/external/article/316110/vmware-cloud-foundation-nsxt-edge-clust.html>.

Failure when deploying multiple isolated workload domains with the same SSO domain in parallel

If you are deploying more than one isolated workload domain at the same time and those workload domains use the same SSO domain, then only the first workload domain is created successfully. Creation of the additional workload domains fails during validation with a message saying that the SSO domain name is already allocated.

Workaround: Deploy the workload domains sequentially. Wait until the first workload domain deploys successfully and then create the additional workload domains.

Heterogeneous operations "Cluster Creation" and "VI Creation" are not supported to be run in parallel when they are operating against same shared NSX instance.

If there is a running VI Creation workflow operating on an NSX resource, then creating a cluster on domains that are sharing that NSX is not possible.

Workaround: None. The VI Creation workflow should complete before the cluster creation workflow can be started.

Adding host fails when host is on a different VLAN

A host add operation can sometimes fail if the host is on a different VLAN.

1. Before adding the host, add a new portgroup to the VDS for that cluster.
2. Tag the new portgroup with the VLAN ID of the host to be added.
3. Add the Host. This workflow fails at the "Migrate host vmknics to dvs" operation.
4. Locate the failed host in vCenter, and migrate the vmk0 of the host to the new portgroup you created in step 1. For more information, see [Migrate VMkernel Adapters to a vSphere Distributed Switch](#) in the vSphere product documentation.
5. Retry the Add Host operation.

NOTE: If you later remove this host in the future, you must manually remove the portgroup as well if it is not being used by any other host.

Deploying partner services on an NSX workload domain displays an error

Deploying partner services, such as McAfee or Trend, on a workload domain enabled for vSphere Update Manager (VUM), displays the "Configure NSX at cluster level to deploy Service VM" error.

Attach the Transport node profile to the cluster and try deploying the partner service. After the service is deployed, detach the transport node profile from the cluster.

If the witness ESXi version does not match with the host ESXi version in the cluster, vSAN cluster partition may occur

vSAN stretch cluster workflow does not check the ESXi version of the witness host. If the witness ESXi version does not match the host version in the cluster, then vSAN cluster partition may happen.

1. Upgrade the witness host manually with the matching ESXi version using the vCenter VUM functionality.
2. Replace or deploy the witness appliance matching with the ESXi version.

vSAN partition and critical alerts are generated when the witness MTU is not set to 9000

If the MTU of the witness switch in the witness appliance is not set to 9000, the vSAN stretch cluster partition may occur. Set the MTU of the witness switch in the witness appliance to 9000 MTU.

Adding a host to a cluster configured with vLCM images fails if the workload domain is using the Dell Hardware Support Manager (OMIVV)

When you try to add a host to a vSphere cluster that uses vSphere Lifecycle Manager (vLCM) images, the task fails and the domain manager log reports "The host (host-name) is currently not managed by OMIVV." The domain manager logs are located at `/var/log/vmware/vcf/domainmanager` on the SDDC Manager VM.

Update the hosts inventory in OMIVV and retry the add host task in the SDDC Manager UI. See the Dell documentation for information about updating the hosts inventory in OMIVV.

The vSAN Performance Service is not enabled for vSAN clusters when CEIP is not enabled

If you do not enable the VMware Customer Experience Improvement Program (CEIP) in SDDC Manager, when you create a workload domain or add a vSphere cluster to a workload domain, the vSAN Performance Service is not enabled for vSAN clusters. When CEIP is enabled, data from the vSAN Performance Service is provided to VMware and this data is used to aid VMware Support with troubleshooting and for products such as VMware Skyline, a proactive cloud monitoring service. See [Customer Experience Improvement Program](#) for more information on the data collected by CEIP.

Enable CEIP in SDDC Manager. See the [VMware Cloud Foundation Documentation](#). After CEIP is enabled, a scheduled task that enables the vSAN Performance Service on existing clusters in workload domains runs every three hours. The service is also enabled for new workload domains and clusters. To enable the vSAN Performance Service immediately, see the [VMware vSphere Documentation](#).

Creation or expansion of a vSAN cluster with more than 32 hosts fails

By default, a vSAN cluster can grow up to 32 hosts. With large cluster support enabled, a vSAN cluster can grow up to a maximum of 64 hosts. However, even with large cluster support enabled, a creation or expansion task can fail on the subtask **Enable vSAN on vSphere Cluster**.

1. Enable Large Cluster Support for the vSAN cluster in the vSphere Client. If it is already enabled skip to step 2.
 - a. Select the vSAN cluster in the vSphere Client.
 - b. Select **Configure > vSAN > Advanced Options**.
 - c. Enable Large Cluster Support.
 - d. Click **Apply**.
 - e. Click **Yes**.
2. Run a vSAN health check to see which hosts require rebooting.
3. Put the hosts into Maintenance Mode and reboot the hosts.

For more information about large cluster support, see <https://kb.vmware.com/kb/2110081>.

Removing a host from a cluster, deleting a cluster from a workload domain, or deleting a workload domain fails if Service VMs (SVMs) are present

If you deployed an endpoint protection service (such as guest introspection) to a cluster through NSX Data Center, then removing a host from the cluster, deleting the cluster, or deleting the workload domain containing the cluster will fail on the subtask **Enter Maintenance Mode on ESXi Hosts**.

- For host removal: Delete the Service VM from the host and retry the operation.
- For cluster deletion: Delete the service deployment for the cluster and retry the operation.

- For workload domain deletion: Delete the service deployment for all clusters in the workload domain and retry the operation.

vCenter Server overwrites the NFS datastore name when adding a cluster to a VI workload domain

If you add an NFS datastore with the same NFS server IP address, but a different NFS datastore name, as an NFS datastore that already exists in the workload domain, then vCenter Server applies the existing datastore name to the new datastore.

If you want to add an NFS datastore with a different datastore name, then it must use a different NFS server IP address.

VMware Cloud Foundation 5.2.1 on Dell VxRail Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Available Languages](#)
- [Deprecation Notices](#)
- [VMware Cloud Foundation Bill of Materials \(BOM\)](#)
- [Documentation](#)
- [Installation and Upgrade Information](#)
- [VMware Cloud Foundation 5.2.1.1 Release Information](#)
- [Resolved Issues](#)
- [Known Issues](#)

Introduction

VMware Cloud Foundation 5.2.1 | 09 OCT 2024 | Build 24307856

VMware Cloud Foundation 5.2.1.1 | 05 DEC 2024 | Build 24397777

Check for additions and updates to these release notes.

What's New

The VMware Cloud Foundation (VCF) 5.2.1 release includes the following:

- **Reduced Downtime Upgrade (RDU) support for vCenter:** VCF users can now leverage vCenter Reduced Downtime Upgrade (RDU) to execute a vCenter upgrade. vCenter RDU is a migration-based approach to upgrading vCenter and reduces the vCenter downtime to less than 5 minutes.
- **NSX in-place upgrades for clusters that use vSphere Lifecycle Manager baselines:** VCF users now have the choice to perform NSX in-place upgrade for clusters that use vSphere Lifecycle Manager baselines. In-place upgrades eliminate the need to place hosts into maintenance mode during the upgrade.
- **Support for the "License Now" option for vSAN add-on licenses based on capacity per tebibyte (TiB):** VCF users can now apply the vSAN TiB capacity license within the SDDC Manager UI to expand storage capacity for their workload domains and clusters. You can also use the "License Later" option to assign the per-TiB vSAN license key using the vSphere Client.
- **Manage all SDDC certificates and passwords from a single UI:** SDDC Manager certificate and password management functionality is now integrated in the vSphere Client to simplify and speed-up day-to-day operations. VCF users can now manage the certificates, integrated certificate authorities, and system user passwords from the Administration section in the vSphere Client.

Available Languages

Beginning with the next major release, VCF will be supporting the following localization languages:

- Japanese
- Spanish
- French

The following languages will no longer be supported:

- Italian, German, and Simplified Chinese.

Impact:

- Customers who have been using the deprecated languages will no longer receive updates in these languages.
- All user interfaces and help documentation will be available only in English or in the three supported languages mentioned above.

Because VCF localization utilizes the browser language settings, ensure that your settings match the desired language.

Deprecation Notices

In a future major release, the following APIs will not be supported:

- POST /v1/bundles
- POST /v1/product-version-catalog

The APIs will be replaced with:

- POST /v1/product-binaries
- POST /v1/product-version-catalogs

After upgrading SDDC Manager to this new release, automation customers who use the unsupported APIs should transition to the new APIs.

- The following features are being deprecated and will be removed in a future major release:
 - Cloud Builder Appliance
 - Cloud Builder APIs
 - Cloud Builder deployment parameter workbooks
 - NSX Edge management workflow
- VMware End Of Availability of Perpetual Licensing and SaaS Services. See <https://blogs.vmware.com/cloud-foundation/2024/01/22/vmware-end-of-availability-of-perpetual-licensing-and-saas-services/> for more information.
- In a future release, the "Connect Workload Domains" option from the VMware Aria Operations card located in **SDDC Manager > Administration > Aria Suite** section will be removed and related VCF Public API options will be deprecated.

Starting with VMware Aria Operations 8.10, functionality for connecting VCF Workload Domains to VMware Aria Operations is available directly from the UI. Users are encouraged to use this method within the VMware Aria Operations UI for connecting VCF workload domains, even if the integration was originally set up using SDDC Manager.

- Deprecation announcements for VMware NSX. See the [VMware NSX 4.2.1 Release Notes](#) for details.

VMware Cloud Foundation Bill of Materials (BOM)

The VMware Cloud Foundation software product is comprised of the following software Bill-of-Materials (BOM). The components in the BOM are interoperable and compatible.

Software Component	Version	Date	Build Number
Cloud Builder VM	5.2.1	09 OCT 2024	24307856
SDDC Manager	5.2.1	09 OCT 2024	24307856

Table continued on next page

Continued from previous page

Software Component	Version	Date	Build Number
VxRail Manager	8.0.310	17 OCT 2024	N/A
VMware vCenter Server Appliance	8.0 Update 3c	09 OCT 2024	24305161
VMware vSAN Witness Appliance	8.0 Update 3	19 JUN 2024	24022510
VMware NSX	4.2.1	09 OCT 2024	24304122
VMware Aria Suite Lifecycle	8.18	23 JUL 2024	24029603

- VMware ESXi and VMware vSAN are part of the VxRail BOM.
- You can use VMware Aria Suite Lifecycle to deploy VMware Aria Automation, VMware Aria Operations, VMware Aria Operations for Logs, and Workspace ONE Access (formerly known as VMware Identity Manager). VMware Aria Suite Lifecycle determines which versions of these products are compatible and only allows you to install/upgrade to supported versions.
- VMware Aria Operations for Logs content packs are installed when you deploy VMware Aria Operations for Logs.
- The VMware Aria Operations management pack is installed when you deploy VMware Aria Operations.
- You can access the latest versions of the content packs for VMware Aria Operations for Logs from the VMware Solution Exchange and the VMware Aria Operations for Logs in-product marketplace store.

Documentation

The following documentation is available:

- [VMware Cloud Foundation on Dell VxRail Guide](#)
- [VMware Cloud Foundation 5.2.1 Release Notes](#)
- [Support Matrix of VMware Cloud Foundation on Dell VxRail](#)

Installation and Upgrade Information

You can perform a sequential or skip level upgrade to VMware Cloud Foundation 5.2.1 on Dell VxRail from VMware Cloud Foundation 4.5 or later. If your environment is at a version earlier than 4.5, you must upgrade the management domain and all VI workload domains to VMware Cloud Foundation 4.5 and then upgrade to VMware Cloud Foundation 5.2.1.

IMPORTANT:

- Before you upgrade a vCenter Server, take a file-based backup. See [Manually Back Up vCenter Server](#).

NOTE: Scripts that rely on SSH being activated on ESXi hosts will not work after upgrading to VMware Cloud Foundation 5.2.1, since VMware Cloud Foundation 5.2.1 deactivates the SSH service by default. Update your scripts to account for this new behavior. See [KB 86230](#) for information about activating and deactivating the SSH service on ESXi hosts.

VMware Cloud Foundation 5.2.1.1 Release Information

VMware Cloud Foundation 5.2.1.1 includes bug fixes, as well as support for stretching vSAN ESA clusters.

You can upgrade to VMware Cloud Foundation 5.2.1.1 from VMware Cloud Foundation 5.2 or later.

Upgrading from 5.2.1:

See [Patching the Management and Workload Domains](#). When selecting the target versions for SDDC Manager and VxRail Manager, choose the versions listed in the BOM update table below.

Upgrading from 5.2:

See [Flexible BOM Upgrade in VMware Cloud Foundation](#). When selecting the target versions for SDDC Manager and VxRail Manager, choose the versions listed in the BOM update table below.

Known issues:

- In order to upgrade from 5.2 to 5.2.1.1, you must download the bundles for both SDDC Manager 5.2.1.0 and SDDC Manager 5.2.1.1.
- The Bundle Management window in the SDDC Manager UI displays "VMware Cloud Foundation Update 5.2.1.0" instead of "VMware Cloud Foundation Update 5.2.1.1" for the 5.2.1.1 bundle. The description of the bundle correctly describes it as the upgrade bundle for 5.2.1.1. This is a cosmetic issue only and does not impact the upgrade.

VMware Cloud Foundation 5.2.1.1 contains the following BOM updates:

Software Component	Version	Date	Build Number
SDDC Manager	5.2.1.1	05 DEC 2024	24397777
VxRail Manager	8.0.311	05 DEC 2024	28839943

Resolved Issues

The following issues have been resolved:

- **Once VCF has been upgraded to 5.2.1, using the Bundle Transfer Utility to upload VxRail patches fails**

Known Issues

For VMware Cloud Foundation 5.2.1 known issues, see [VMware Cloud Foundation 5.2.1 known issues](#). Some of the known issues may be for features that are not available on VMware Cloud Foundation on Dell VxRail.

VMware Cloud Foundation 5.2.1 on Dell VxRail Known Issues appear below:

Add host validation from the VxRail vCenter plugin fails with the error "Error while preparing network for VLAN validation on VMware VDS"

When you use the VxRail vCenter plugin to add a host to a cluster in a workload domain where the vSphere distributed switch (VDS) for the Management network has its load balancing policy configured as "Route based on IP hash", host validation fails.

Workaround: See <https://www.dell.com/support/kbdoc/000236639>.

SDDC Manager fails on host discovery

VxRail host discovery fails in the SDDC Manager UI with the following error: "Failed to load Host cluster details".

Workaround: Upgrade to VMware Cloud Foundation 5.2.1.1, which resolves the issue, or add the VxRail host(s) via the API Explorer. See <https://knowledge.broadcom.com/external/article?articleNumber=379202>.

Error while retrying failed VxRail upgrade to 5.2.1

Retrying a failed VxRail upgrade fails with the error message `Failed to process the request body`.

Workaround: See <https://www.dell.com/support/kbdoc/000227705>.

Compatibility warning when upgrading from VMware Cloud Foundation 5.1 to 5.2.1

When you plan an upgrade from VMware Cloud Foundation 5.1 to VMware Cloud Foundation 5.2.1 a compatibility warning may display.

VMware Cloud Foundation 5.2.0.0 with VxRail Manager 8.0.300

Select Version

Unable to verify the compatibility for the following product versions. Please check the product documentation before proceeding to upgrade: .

Workaround: This upgrade path is fully-supported and you can ignore the warning.

Error during thumbprint verification when adding hosts to a cluster

When you add hosts to a VxRail cluster that is part of a workload domain that uses vLCM images, you may see the following error: Dependency on Enabling SSH could not be performed on one or many host in the request payload.

Workaround: Enable SSH on the host(s) and retry the task.

VxRail first run fails with error "Failed to do vLCM remediation for cluster"

If an NVMe device or controllers are not VMware Certified, the VxRail first run fails with a Skyline Health alert.

Workaround: Browse to the Skyline Health section for the cluster using the vSphere Client, silence the alert, and retry the first run.

Creating a VI workload domain or adding a cluster with VMFS on FC storage fails

When you use the SDDC Manager UI to create a VI workload domain or add a cluster with VMFS on FC storage, the task fails.

Workaround: Use the Workflow Optimization script or VMware Cloud Foundation on Dell VxRail API to create the VI workload domain or add the cluster.

VxRail plugin missing after the vCenter Server certificate or password is rotated from SDDC Manager

When you rotate a vCenter Server certificate or password, the VxRail plugin may disappear for clusters managed by another vCenter Server.

Workaround: See <https://www.dell.com/support/kbdoc/000224478>.

Creating a workload domain or adding a cluster using ESXi hosts with GPU drivers

If your ESXi hosts include a GPU driver, you must upload the GPU driver to the VxRail Manager before you:

- Create a VI workload domain that uses vSphere Lifecycle Manager images.
- Add a cluster to a VI workload domain that uses vSphere Lifecycle Manager images.

Workaround: See <https://www.dell.com/support/kbdoc/en-in/000202491>.

VxRail does not allow more than 2 active/one active uplink, 1 standby nics.

Mapping more than 2 uplinks causes VxRail cluster validation to fail during the VxRail Dry Run.

Workaround: When creating a custom NIC profile, map no more than 2 uplinks to active/standby uplinks.

Support for consuming VxRail upgrade bundles for 5x-5y is unavailable

VxRail Manager 8.0.x to 8.0.300 upgrade fails stating a 8.0.x.zip file does not exist error.

Workaround: See KB article [94747](#) for the scripts and manual steps to mitigate this compatibility gap.

Failed VxRail first run prevents new cluster/workload domain creation

If the VxRail first run fails, some objects associated with the failed task remain in the vCenter Server inventory and prevent new cluster/workload domain creation.

Workaround: Remove the inventory objects associated with the failed task using the vSphere Client.

1. Log in to the vSphere Client.
2. In the Hosts and Clusters inventory, right-click the failed cluster and select **vSAN > Shutdown cluster**.

3. After the shutdown completes, right-click the failed cluster and select **Delete**.

After the inventory is cleaned up, you can retry adding a cluster or creating a workload domain.

Add VxRail hosts validation fails

When adding VxRail hosts to a workload domain or cluster that uses Fibre Channel (FC) storage, the task may fail with the message No shared datastore can be found on host. This can happen if you used the Workflow Optimization Script to deploy the workload domain or cluster and chose an FC datastore name other than the default name.

Workaround: Use the VMware Host Client to rename the FC datastore on the new VxRail hosts to match the name you entered when creating the workload domain or cluster. Once the FC datastore name of the new hosts matches the existing FC datastore name, retry the Add VxRail Hosts operation.

Unsupported versions of VxRail not restricted during create cluster/domain operations

VCF on VxRail does not restrict using unsupported/unpaired versions of VxRail for create cluster/domain operations. If nodes are re-imaged with a VxRail version that is not paired with the current VCF release, VCF does not restrict using these nodes for creating a cluster/domain.

Workaround: Use the VxRail Manager version paired with the correct VCF release for create domain/cluster operations.

vSAN/vMotion network disruption can occur when using the workflow optimization script

When you use the workflow optimization script to create a new VI workload domain or add a new cluster to an existing workload domain, you can cause a network disruption on existing vSAN/vMotion networks if:

- The IP range for the new vSAN network overlaps with the IP range for an existing vSAN network.
- The IP range for the new vMotion network overlaps with the IP range for an existing vMotion network.

Workaround: None. Make sure to provide vSAN/vMotion IP ranges that do not overlap with existing vSAN/vMotion networks.

VMware Cloud Foundation 5.2 Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Available Languages](#)
- [Deprecation Notices](#)
- [VMware Cloud Foundation Bill of Materials \(BOM\)](#)
- [Supported Hardware](#)
- [Documentation](#)
- [Browser Compatibility and Screen Resolutions](#)
- [Installation and Upgrade Information](#)
- [Resolved Issues](#)
- [Known Issues](#)

Introduction

VMware Cloud Foundation 5.2 | 23 JUL 2024 | Build 24108943

Check for additions and updates to these release notes.

What's New

The VMware Cloud Foundation (VCF) 5.2 release includes the following:

- **Support for Identity Federation with Entra ID:** VCF users can now configure Microsoft Entra ID (formerly known as Azure AD) as an identity provider.
- **APIs for Auditing PCI Compliance:** VCF users can now use a new set of APIs that audit VCF configuration for compliance with 9 relevant PCI-DSS controls.
- **vSAN Max support:** vSAN Max is a disaggregated storage offering which enables petabyte scale storage-only clusters. vSAN Max is powered by ESA as the underlying storage platform, which is a high-performance file system that can scale up to high densities with no penalty to performance. ESA also provides other benefits such as built-in, efficient, scalable snapshots, and low overhead data services such as encryption and compression.
- **vSAN ESA Stretched Cluster:** VCF users can now configure ESA Stretched Cluster in vSAN Ready Nodes. It enables customers to take the concept of fault domains to protect an environment spanning two physical sites from downtime in the event of a site failure.
- **VCF Import Tool (for vSphere & vSAN):** The VCF Import Tool integrates existing vSphere environments into VMware Cloud Foundation, centralizing management and optimizing resources without needing a full rebuild.
- **Support for additional principal storage types with a converted management domain:** If you use the VCF Import Tool to convert an existing vSphere environment to a VCF management domain, that management domain can use VMFS-FC and NFS v3 as principal storage, in addition to vSAN.
- **Dual DPU Support:** VCF users can now leverage Dual DPU support. Dual DPU support boosts availability and performance. Active/Standby ensures continuity against failures, while dual independent DPUs double offload capacity and provide isolation.
- **Avi Load Balancer Integration with VCF:** VCF users can now deploy Avi (formerly NSX Advanced Load Balancer) as part of a new workload domain and perform password rotation and certificate management of the ALB infrastructure from SDDC Manager.
- **Out of Band Changes from vCenter:** Out of Band changes from vCenter can be manually synced with SDDC Manager. This includes inventory changes (for example, adding a host to a cluster) and object name changes (for example, datacenter name, datastore name, port group name).
- **ESXi Live Patching :** VCF users can now apply ESXi security patches without requiring VM evacuation on ESXi hosts.
- **Flexible Target BOM for Upgrades:** VCF users can now create a composite and customized BOM using patches when upgrading workload domains. Customers can plan an upgrade along with patches in one orchestrated workflow instead of performing an upgrade and applying patches in separate maintenance windows.
- **Async Patching with SDDC Manager:** Customers previously used the standalone Async Patch Tool to apply patches to the VCF BOM components. VCF 5.2 provides the ability to apply BOM component patches from the SDDC Manager UI.
- **Day N workflows with Embedded Async Patching:** VCF users can now add new workload domains and clusters with patched versions of individual BOM components from SDDC Manager.
- **Asynchronous SDDC Manager Upgrades:** VCF users can now upgrade SDDC Manager independently from the rest of the BOM to apply critical fixes, security patches, and to enable specific features related to SDDC Manager.
- **Authenticated Proxy:** VCF users can now use proxy authentication from SDDC Manager to enable secure connectivity from SDDC Manager to the internet.
- **Offline Depot:** VCF users can now perform lifecycle bundle downloads in offline/air-gapped environments in a simplified manner. The offline depot downloads and stages VCF SDDC Manager and BOM component bundles and enables customers to configure SDDC Manager to download the bundles directly from the offline depot.
- **Isolated Workload Domains Sharing NSX:** VCF users can now create and manage isolated workload domains that can share an NSX Manager instance between them.

Available Languages

Beginning with the next major release, VCF will be supporting the following localization languages:

- Japanese
- Spanish
- French

The following languages will no longer be supported:

- Italian, German, and Simplified Chinese.

Impact:

- Customers who have been using the deprecated languages will no longer receive updates in these languages.
- All user interfaces and help documentation will be available only in English or in the three supported languages mentioned above.

Because VCF localization utilizes the browser language settings, ensure that your settings match the desired language.

Deprecation Notices

- VMware End Of Availability of Perpetual Licensing and SaaS Services. See <https://blogs.vmware.com/cloud-foundation/2024/01/22/vmware-end-of-availability-of-perpetual-licensing-and-saas-services/> for more information.
- The Composable Infrastructure feature is deprecated and removed.
- In a future release, the "Connect Workload Domains" option from the VMware Aria Operations card located in **SDDC Manager > Administration > Aria Suite** section will be removed and related VCF Public API options will be deprecated.

Starting with VMware Aria Operations 8.10, functionality for connecting VCF Workload Domains to VMware Aria Operations is available directly from the UI. Users are encouraged to use this method within the VMware Aria Operations UI for connecting VCF workload domains, even if the integration was originally set up using SDDC Manager.

- Deprecation announcements for VMware NSX. See the [VMware NSX 4.2.0 Release Notes](#) for details.
 - NSX Manager APIs and NSX Advanced UIs
 - NSX embedded (NSXe)
 - Some parameters in Switch IPFIX
 - NSX Migration Coordinator

VMware Cloud Foundation Bill of Materials (BOM)

The VMware Cloud Foundation software product is comprised of the following software Bill-of-Materials (BOM). The components in the BOM are interoperable and compatible.

Software Component	Version	Date	Build Number
Cloud Builder VM	5.2	23 JUL 2024	24108943
SDDC Manager	5.2	23 JUL 2024	24108943
VMware vCenter Server Appliance	8.0 Update 3a	18 JUL 2024	24091160
VMware ESXi	8.0 Update 3	25 JUN 2024	24022510
VMware vSAN Witness Appliance	8.0 Update 3	25 JUN 2024	24022510
VMware NSX	4.2	23 JUL 2024	24105817
VMware Aria Suite Lifecycle	8.18	23 JUL 2024	24029603

- VMware vSAN is included in the VMware ESXi bundle.
- You can use VMware Aria Suite Lifecycle to deploy VMware Aria Automation, VMware Aria Operations, VMware Aria Operations for Logs, and Workspace ONE Access. VMware Aria Suite Lifecycle determines which versions of these products are compatible and only allows you to install/upgrade to supported versions.
- VMware Aria Operations for Logs content packs are installed when you deploy VMware Aria Operations for Logs.
- The VMware Aria Operations management pack is installed when you deploy VMware Aria Operations.
- You can access the latest versions of the content packs for VMware Aria Operations for Logs from the VMware Solution Exchange and the VMware Aria Operations for Logs in-product marketplace store.

Supported Hardware

For details on supported configurations, see the [VMware Compatibility Guide \(VCG\)](#) and the Hardware Requirements section on the Prerequisite Checklist tab in the [Planning and Preparation Workbook](#).

Documentation

To access the VCF documentation, go to the [VMware Cloud Foundation product documentation](#).

To access the documentation for VMware software products that SDDC Manager can deploy, see the product documentation and use the drop-down menus on the page to choose the appropriate version:

- [VMware vSphere product documentation](#), includes the documentation for ESXi and vCenter Server
- [VMware vSAN product documentation](#)
- [VMware NSX product documentation](#)

Browser Compatibility and Screen Resolutions

The VMware Cloud Foundation web-based interface supports the latest two versions of the following web browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge

For the Web-based user interfaces, the supported standard resolution is 1920 by 1080 pixels.

Installation and Upgrade Information

You can install VMware Cloud Foundation 5.2 as a new release or perform a sequential or skip-level upgrade to VMware Cloud Foundation 5.2.

Installing as a New Release

The new installation process has three phases:

- **Phase One: Prepare the Environment:** The [Planning and Preparation Workbook](#) provides detailed information about the software, tools, and external services that are required to implement a Software-Defined Data Center (SDDC) with VMware Cloud Foundation, using a standard architecture model.
- **Phase Two: Image all servers with ESXi:** Image all servers with the ESXi version mentioned in the Cloud Foundation Bill of Materials (BOM) section. See the [VMware Cloud Foundation Deployment Guide](#) for information on installing ESXi.
- **Phase Three: Install Cloud Foundation 5.2:** See the [VMware Cloud Foundation Deployment Guide](#) for information on deploying Cloud Foundation.

Upgrading to Cloud Foundation 5.2

You can perform a sequential or skip-level upgrade to VMware Cloud Foundation 5.2 from VMware Cloud Foundation 4.5.0 or later. If your environment is at a version earlier than 4.5.0, you must upgrade the management domain and all VI workload domains to VMware Cloud Foundation 4.5.0 or above and then upgrade to VMware Cloud Foundation 5.2. For more information see [VMware Cloud Foundation Lifecycle Management](#).

IMPORTANT

Before you upgrade a vCenter Server, take a file-based backup. See [Manually Back Up vCenter Server](#).

NOTE

Since VMware Cloud Foundation disables the SSH service by default, scripts that rely on SSH being enabled on ESXi hosts will not work after upgrading to VMware Cloud Foundation 5.2. Update your scripts to account for this new behavior. See [KB 86230](#) for information about enabling and disabling the SSH service on ESXi hosts.

Resolved Issues

The following issue is resolved in this release:

- VCF ESXi Upgrade with 'quick boot' option fails for hosts configured with TPM
- Deploying the management domain using vSphere Lifecycle Manager (vLCM) images fails
- VM MANAGEMENT port group may get created on the wrong vSphere Distributed Switch (VDS)
- Entering pNICs in non-lexicographic order in the deployment parameter workbook does not work as expected
- Entering more than two pNICs for the primary vDS in the deployment parameter workbook does not work as expected
- SDDC Manager UI is showing a vSAN License as Active even though it was not assigned
- Filtering bundles by "Downloaded" in the SDDC Manager UI does not show any results
- Cannot add unused vmnics to an existing vSphere Distributed Switch (VDS)

Known Issues

VMware Cloud Foundation Known Issues

VCF Import Tool does not support clusters that use vSphere Configuration Profiles

If you use the VCF Import Tool to import/convert an existing vSphere environment that includes clusters that use vSphere Configuration Profiles, the task fails during NSX deployment.

Workaround: None. Clusters that use vSphere Configuration Profiles do not support NSX.

VMware Cloud Foundation 5.2 does not support the "License Now" option for vSAN add-on licenses based on capacity per tebibyte (TiB)

With vSphere 8.0 Update 3, you can license your use of vSAN storage based on TiB capacity. When using a per-TiB vSAN license with VMware Cloud Foundation, if you enter the license key using the "License Now" option either during bringup of the management domain, or when deploying or expanding a VI workload domain, the workflow fails.

Workaround: Use the "License Later" option and assign the per-TiB vSAN license key later using the vSphere Client.

Primary datastore is not getting set for imported workload domains with NFS 4.1 datastore

When you use the VCF Import Tool to import a cluster for which NFS 4.1 is the only shared datastore, the primary datastore and datastore type is not getting set in VCF and the workload domain is not visible in the SDDC Manager UI. See <https://knowledge.broadcom.com/external/article/372424> for details.

Workaround: None.

Limitations for importing vSAN clusters

When you use the VCF Import Tool to import a vSAN cluster, you should avoid importing clusters with certain configurations. SDDC Manager day-N operations will not be supported on imported vSAN clusters with these configurations. See <https://knowledge.broadcom.com/external/article/371494> for details.

Workaround: None.

Lifecycle Management Precheck does not throw an error when NSX Manager inventory is out of sync

The Lifecycle Management Precheck displays a green status and does not generate any errors for NSX Manager inventory.

Workaround: None

Upgrade Pre-Check Scope dropdown may contain additional entries

When performing Upgrade Prechecks through SDDC Manager UI and selecting a target VCF version, the Pre-Check Scope dropdown may contain more selectable entries than necessary. SDDC Manager may appear as an entry more than once. It also may be included as a selectable component for VI domains, although it's a component of the management domain.

Workaround: None. The issue is visual with no functional impact.

Converting clusters from vSphere Lifecycle Manager baselines to vSphere Lifecycle Manager images is not supported.

vSphere Lifecycle Manager baselines (previously known as vSphere Update Manager or VUM) are deprecated in vSphere 8.0, but continue to be supported. See [KB article 89519](#) for more information.

VMware Cloud Foundation 5.0 does not support converting clusters from vSphere Lifecycle Manager baselines to vSphere Lifecycle Manager images. This capability will be supported in a future release.

Workaround: None

Workload Management and NSX Federation

While you cannot deploy Workload Management (vSphere with Tanzu) to a workload domain using stretched NSX segments and T1/T0 when that workload domain's NSX Data Center instance is participating in an NSX Federation, you can deploy NSX local segments and local dedicated T0/T1 from NSX Local Manager that are not stretched between two VCF instances using NSX Federation. Make sure to configure NSX Federation before deploying Workload Management to avoid any potential NSX import issues to NSX Global Manager.

Workaround: None.

Upgrade Known Issues

Bundle Transfer Utility fails to upload the NSX Advanced Load Balancer install bundle

If you are on a pre-5.2 version of VMware Cloud Foundation and use the Bundle Transfer Utility to download all bundles for VCF 5.2, then uploading the NSX Advanced Load Balancer install bundle fails. This bundle is only supported with SDDC Manager 5.2 and later.

Workaround: Upgrade SDDC Manager to 5.2 and then retry uploading the NSX Advanced Load Balancer install bundle.

NSX host cluster upgrade fails

If you are upgrading a workload domain that uses vSphere Lifecycle Manager images and its cluster image was created from an ESXi host that uses vSphere Lifecycle Manager baselines, then NSX host cluster upgrade will fail. A cluster image created from an ESXi host that uses vSphere Lifecycle Manager baselines contains an NSX component that causes this issue.

NOTE: This issue is resolved if you have ESXi and vCenter Server 8.0 Update 3 or later.

Workaround: Do not create cluster images from an ESXi host that uses vSphere Lifecycle Manager baselines. If you encounter this issue, you can resolve it by using the vSphere Client to remove the NSX LCP Bundle component from the cluster image.

The screenshot shows the 'Edit Image' configuration page in the vSphere Client. The 'Components' section is expanded, showing a table with the following data:

Component Name	Version	Notes
NSX LCP Bundle	NSX LCP Bundle(4.1.2.0.0-8.0.22305537)	Manually added component

A red arrow points to a trash icon in the 'Notes' column next to the 'NSX LCP Bundle' component, indicating it should be removed.

SDDC Manager UI shows the incorrect source version when upgrading SDDC Manager

When you view the VMware Cloud Foundation Update Status for SDDC Manager, the UI may show the incorrect source version.



Workaround: None. This is a cosmetic issue only and does not affect the upgrade.

Workspace ONE Access inventory sync fails in SDDC Manager after upgrading VMware Aria Suite Lifecycle

After upgrading Aria Suite Lifecycle to version 8.12 or later, triggering a Workspace ONE Access inventory sync from Aria Suite Lifecycle fails. The SDDC Manager UI reports the following error: Failed to configure WSA <wsa_fqdn> in vROps .vrops_fqdn>, because Failed to manage vROps adapter.

Workaround: Download the bundle for your version of Aria Suite Lifecycle to SDDC Manager and retry the inventory sync.

VCF ESXi upgrade fails during post validation due to HA related cluster configuration issue

The upgrade of ESXi Cluster fails with error that is similar to below error message:

```
Cluster Configuration Issue: vSphere HA failover operation in progress in cluster
<cluster-name> in datacenter <datacenter-name>: 0 VMs being restarted, 1 VMs waiting for a
retry, 0 VMs waiting for resources, 0 inaccessible vSAN VMs
```

Workaround: See [KB article 90985](#).

Lifecycle Management Precheck does not throw an error when NSX Manager inventory is out of sync

Workaround None.

NSX upgrade may fail if there are any active alarms in NSX Manager

If there are any active alarms in NSX Manager, the NSX upgrade may fail.

Workaround: Check the NSX Manager UI for active alarms prior to NSX upgrade and resolve them, if any. If the alarms are not resolved, the NSX upgrade will fail. The upgrade can be retried once the alarms are resolved.

SDDC Manager upgrade fails at "Setup Common Appliance Platform"

If a virtual machine reconfiguration task (for example, removing a snapshot or running a backup) is taking place in the management domain at the same time you are upgrading SDDC Manager, the upgrade may fail.

Workaround: Schedule SDDC Manager upgrades for a time when no virtual machine reconfiguration tasks are happening in the management domain. If you encounter this issue, wait for the other tasks to complete and then retry the upgrade.

Parallel upgrades of vCenter Server are not supported

If you attempt to upgrade vCenter Server for multiple VI workload domains at the same time, the upgrade may fail while changing the permissions for the vpostgres configuration directory in the appliance. The message `chown -R vpostgres:vpgmongrp /storage/archive/vpostgres` appears in the PatchRunner.log file on the vCenter Server Appliance.

Workaround: Each vCenter Server instance must be upgraded separately.

When you upgrade VMware Cloud Foundation, one of the vSphere Cluster Services (vCLS) agent VMs gets placed on local storage

vSphere Cluster Services (vCLS) ensures that cluster services remain available, even when the vCenter Server is unavailable. vCLS deploys three vCLS agent virtual machines to maintain cluster services health. When you upgrade VMware Cloud Foundation, one of the vCLS VMs may get placed on local storage instead of shared storage. This could cause issues if you delete the ESXi host on which the VM is stored.

Workaround: Deactivate and reactivate vCLS on the cluster to deploy all the vCLS agent VMs to shared storage.

1. Check the placement of the vCLS agent VMs for each cluster in your environment.
 - a. In the vSphere Client, select **Menu > VMs and Templates**.
 - b. Expand the vCLS folder.
 - c. Select the first vCLS agent VM and click the Summary tab.
 - d. In the Related Objects section, check the datastore listed for Storage. It should be the vSAN datastore. If a vCLS agent VM is on local storage, you need to deactivate vCLS for the cluster and then re-enable it.
 - e. Repeat these steps for all vCLS agent VMs.

2. Deactivate vCLS for clusters that have vCLS agent VMs on local storage.

- a. In the vSphere Client, click **Menu > Hosts and Clusters**.
- b. Select a cluster that has a vCLS agent VM on local storage.
- c. In the web browser address bar, note the moref id for the cluster.

For example, if the URL displays as `https://vcenter-1.vrack.vsphere.local/ui/app/cluster;nav=urn:vmomi:ClusterComputeResource:domain-c8:503a0d38-442a-446f-b283-d3611bf035fb/summary`, then the moref id is `domain-c8`.

- d. Select the vCenter Server containing the cluster.
- e. Click **Configure > Advanced Settings**.
- f. Click **Edit Settings**.
- g. Change the value for `config.vcls.clusters.<moref id>.enabled` to `false` and click **Save**.

If the `config.vcls.clusters.<moref id>.enabled` setting does not appear for your moref id, then enter its Name and `false` for the Value and click **Add**.

- h. Wait a couple of minutes for the vCLS agent VMs to be powered off and deleted. You can monitor progress in the Recent Tasks pane.

3. Enable vCLS for the cluster to place the vCLS agent VMs on shared storage.

- a. Select the vCenter Server containing the cluster and click **Configure > Advanced Settings**.
- b. Click **Edit Settings**.
- c. Change the value for `config.vcls.clusters.<moref id>.enabled` to `true` and click **Save**.
- d. Wait a couple of minutes for the vCLS agent VMs to be deployed and powered on. You can monitor progress in the Recent Tasks pane.

4. Check the placement of the vCLS agent VMs to make sure they are all on shared storage

You are unable to update NSX Data Center in the management domain or in a workload domain with vSAN principal storage because of an error during the NSX transport node precheck stage

In SDDC Manager, when you run the upgrade precheck before updating NSX Data Center, the NSX transport node validation results with the following error.

No coredump target has been configured. Host core dumps cannot be saved.:System logs on host sfo01-m01-esx04.sfo.rainpole.io are stored on non-persistent storage. Consult product documentation to configure a syslog server or a scratch partition.

Because the upgrade precheck results with an error, you cannot proceed with updating the NSX Data Center instance in the domain. VMware Validated Design supports vSAN as the principal storage in the management domain. However, vSAN datastores do not support scratch partitions. See VMware Knowledge Base article [2074026](#).

Disable the update precheck validation for the subsequent NSX Data Center update.



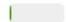

1. Log in to SDDC Manager as **vcf** using a Secure Shell (SSH) client.
2. Open the `application-prod.properties` file for editing: `vi /opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties`
3. Add the following property and save the file: `lcm.nsxt.suppress.prechecks=true`
4. Restart the life cycle management service: `systemctl restart lcm`
5. Log in to the SDDC Manager user interface and proceed with the update of NSX Data Center.

Bring-up Known Issues

SDDC Manager Known Issues

The SDDC Manager UI displays incorrect CPU and memory utilization information for hosts

When you view CPU and memory usage for hosts in the SDDC Manager UI (**Inventory > Hosts**), the information may not reflect actual utilization.

FQDN	Host IP	Network Pool	Configuration Status	Host State	Cluster	CPU Usage	Memory Usage
esxi-1.vrack.vsphere.local	10.0.0.100	bringup-networkpool 	 Active	Assigned (sddcid-1001)	SDDC-Cluster1	4% 	51% 

Workaround: Use the vSphere Client to view CPU and memory utilization information for hosts.

Install bundle for VMware Aria Suite Lifecycle 8.18 displays incomplete information

The SDDC Manager UI displays incomplete information for the VMware Aria Suite Lifecycle 8.18 install bundle.

VMware Software Install Bundle - vRealize Suite Lifecycle Manager
8.18.0-24029603
Released Jun 14, 2024 2 GB
This VMware Software Upgrade contains [Add details here]. For more information, please see - [Insert link to release notes / KB article here]
[View Details](#)

Workaround: None. This is a cosmetic issue and does not impact your ability to download or use the bundle.

When creating a network pool, the IP addresses you provide are not validated to ensure that they are not in use
SDDC Manager *does* validate the included IPs for a new network pool against other network pools and against all other network types (vSAN, NFS, and so on) being added to the new network pool. However, it *does not* validate them against other components that are already deployed in the VMware Cloud Foundation instance (for example, ESXi hosts, NSX Managers, and so on). This can result in duplicate IP address errors or failed workflows.

Workaround: When creating a network pool, do not include any IP addresses that are already in use. If you already created a network pool that includes IP addresses that are used by other components, contact Broadcom Support to resolve the issue.

vSphere Lifecycle Manager images that utilize “removed components” are not supported

Starting with vSphere 8.0 Update 3, you can remove the Host Client and VMware Tools components from a base image, remove unnecessary drivers from vendor add-ons and components, and override existing drivers in a desired image. SDDC Manager does not support this functionality yet for imported or extracted cluster images.

Workaround: None.

Workload Domain Known Issues

Deploying Avi Load Balancer fails

When you deploy Avi Load Balancer (formerly known as NSX Advanced Load Balancer), the deployment may fail with the message `OVA upload to NSX failed`. This can happen if the certificates of the management domain NSX Manager nodes do not include their IP addresses in their Subject Alternative Names (SANs).

Workaround: Generate new CSRs for the management domain NSX Manager nodes, making sure to include IP addresses for the SANs. For example:

Subject Alternative Name ×

Enter the SAN for each resource you selected.

	nsx-mgmt-1.vrack.vsphere.local
SAN(s)	172.16.11.132
Optional	E.g. nsx-mgmt-1.vrack.vsphere.local, 1, etc.

	vip-nsx-mgmt.vrack.vsphere.local
SAN(s)	172.16.11.131
Optional	E.g. vip-nsx-mgmt.vrack.vsphere.local, 1, etc.

◀
▶

Generate the signed certificates using the CSRs and then install the signed certificates in the NSX Manager nodes. See [Managing Certificates in VMware Cloud Foundation](#) for more information.

Once the new certificates are installed, retry deploying Avi Load Balancer.

Switch configuration error when deploying a VI workload domain or adding a cluster to a workload domain with hosts that have two DPUs

If you are using ESXi hosts with two data processing units (DPU) to deploy a new VI workload domain or add a cluster to a workload domain, you may see the following error during switch configuration: `Error in validating Config Profiles`. This can be caused by the presence of a vusb0 network adapter on the hosts.

Workaround: Contact Broadcom Support to remove the vusb0 interface from the SDDC Manager inventory.

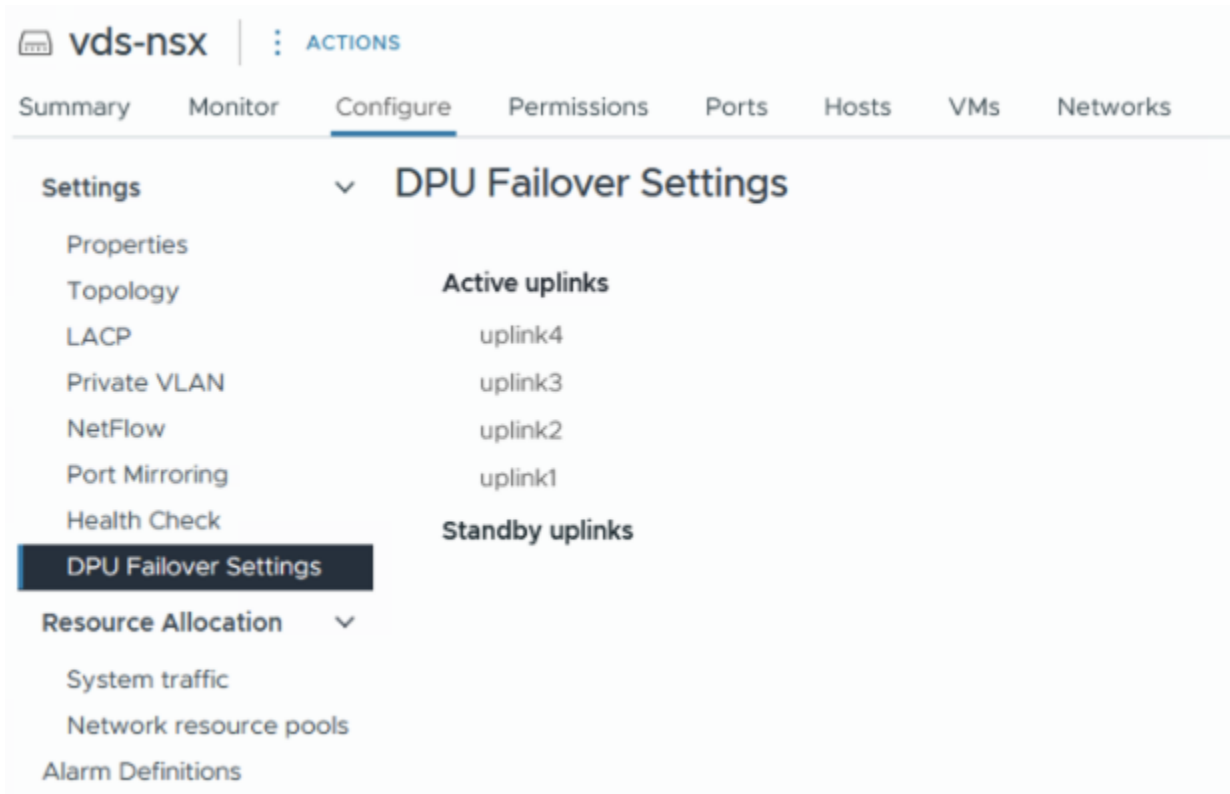
Deploying a VI workload domain or adding a cluster to a workload domain fails with hosts that have two DPUs

If you are using ESXi hosts with two data processing units (DPU) to deploy a new VI workload domain or add a cluster to a workload domain, the task fails when adding the ESXi hosts to the vSphere Distributed Switch (VDS) with the error `Cannot complete a vSphere Distributed Switch operation for one or more host members`.

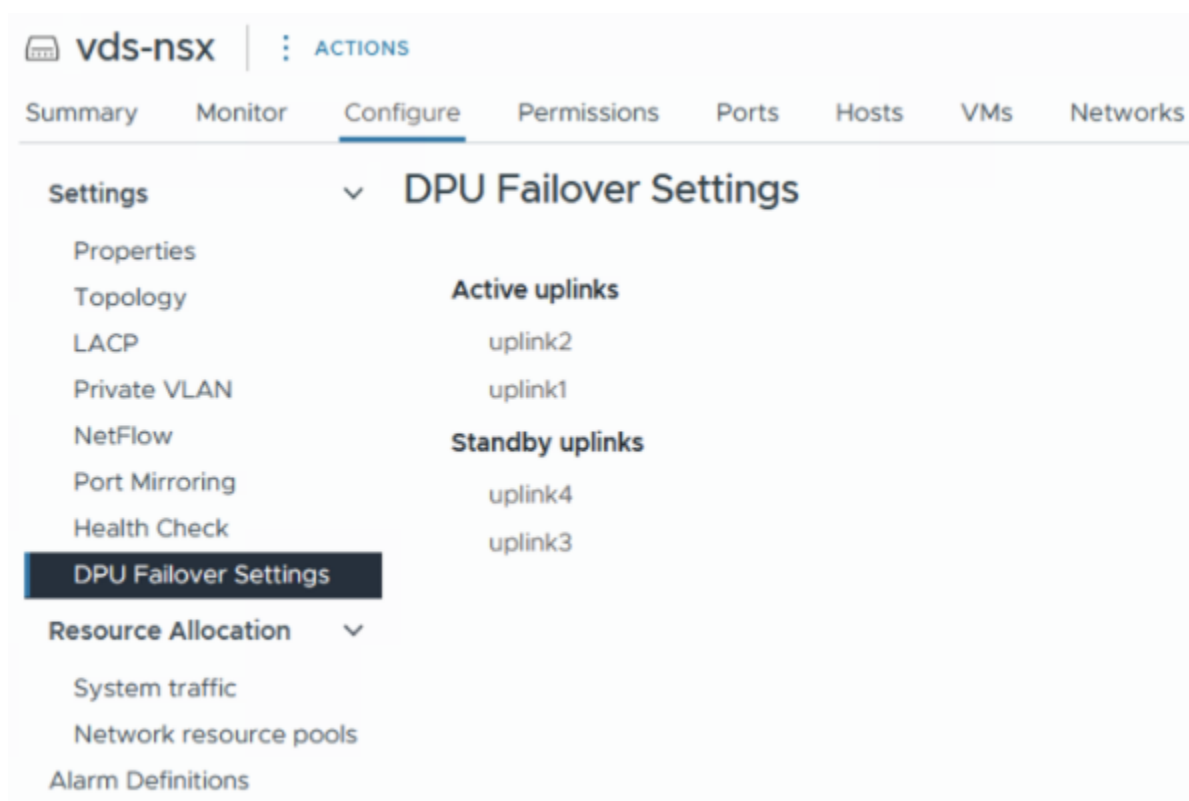
The VDS created by SDDC Manager for dual DPU hosts has all 4 uplinks in Active mode and this does not work with an NSX uplink profile where one set of DPU uplinks is Active and a second set of DPU uplinks is Standby.

Workaround: Use the vSphere Client to manually update the DPU failover settings for the VDS and then retry the workflow from SDDC Manager.

1. In the vSphere Client, browse to the VDS in the vCenter that contains the hosts.
2. Click the **Configure** tab and select **DPU Failover Settings**.



3. Click **Edit** and move uplink3 and uplink4 from Active to Standby.
4. Click **OK**.



5. In the SDDC Manager UI, retry the failed workflow.

NSX Edge cluster deployment fails at "Create VLAN Port Group" stage with message "Invalid parameter: port group already exists"

When you deploy an NSX Edge cluster for VI workload domain and you select the option "USE ESXI MANAGEMENT VMK'S VLAN", the management portgroup name and VLAN ID are auto-populated. SDDC Manager tries to create a portgroup with same VLAN and portgroup name as ESXi management, but since the portgroup name already exists in vCenter the operation fails.

Workaround: If you select the option "USE ESXI MANAGEMENT VMK'S VLAN", change the auto-populated portgroup name to something else so that there is no conflict. If the environment is already in failed state, remove the partially deployed edge cluster. See <https://knowledge.broadcom.com/external/article/316110/vmware-cloud-foundation-nsxt-edge-clust.html>.

Remove unresponsive ESXi Host fails when SDDC Manager certificate does not have subject alternative name

When trying to remove an unresponsive ESXi Host from a cluster, if the SDDC Manager certificate does not have subject alternative name (SAN), the removal of the host will fail at the task "Remove vmknics(s) from ESXi Hosts".

Workaround: Rotate the certificate of the SDDC Manager which will include the SAN. Once the certificate is rotated and it has SAN, the failed remove host workflow can be retried.

Failure when deploying multiple isolated workload domains with the same SSO domain in parallel

If you are deploying more than one isolated workload domain at the same time and those workload domains use the same SSO domain, then only the first workload domain is created successfully. Creation of the additional workload domains fails during validation with a message saying that the SSO domain name is already allocated.

Workaround: Deploy the workload domains sequentially. Wait until the first workload domain deploys successfully and then create the additional workload domains.

Heterogeneous operations "Cluster Creation" and "VI Creation" are not supported to be run in parallel when they are operating against same shared NSX instance.

If there is a running VI Creation workflow operating on an NSX resource, then creating a cluster on domains that are sharing that NSX is not possible.

Workaround: None. The VI Creation workflow should complete before the cluster creation workflow can be started.

Adding host fails when host is on a different VLAN

A host add operation can sometimes fail if the host is on a different VLAN.

1. Before adding the host, add a new portgroup to the VDS for that cluster.
2. Tag the new portgroup with the VLAN ID of the host to be added.
3. Add the Host. This workflow fails at the "Migrate host vmknics to dvs" operation.
4. Locate the failed host in vCenter, and migrate the vmk0 of the host to the new portgroup you created in step 1. For more information, see [Migrate VMkernel Adapters to a vSphere Distributed Switch](#) in the vSphere product documentation.
5. Retry the Add Host operation.

NOTE: If you later remove this host in the future, you must manually remove the portgroup as well if it is not being used by any other host.

Deploying partner services on an NSX workload domain displays an error

Deploying partner services, such as McAfee or Trend, on a workload domain enabled for vSphere Update Manager (VUM), displays the "Configure NSX at cluster level to deploy Service VM" error.

Attach the Transport node profile to the cluster and try deploying the partner service. After the service is deployed, detach the transport node profile from the cluster.

If the witness ESXi version does not match with the host ESXi version in the cluster, vSAN cluster partition may occur

vSAN stretch cluster workflow does not check the ESXi version of the witness host. If the witness ESXi version does not match the host version in the cluster, then vSAN cluster partition may happen.

1. Upgrade the witness host manually with the matching ESXi version using the vCenter VUM functionality.
2. Replace or deploy the witness appliance matching with the ESXi version.

vSAN partition and critical alerts are generated when the witness MTU is not set to 9000

If the MTU of the witness switch in the witness appliance is not set to 9000, the vSAN stretch cluster partition may occur. Set the MTU of the witness switch in the witness appliance to 9000 MTU.

Adding a host to a vLCM-enabled workload domain configured with the Dell Hardware Support Manager (OMIVV) fails

When you try to add a host to a vSphere cluster for a workload domain enabled with vSphere Lifecycle Manager (vLCM), the task fails and the domain manager log reports "The host (host-name) is currently not managed by OMIVV." The domain manager logs are located at `/var/log/vmware/vcf/domainmanager` on the SDDC Manager VM.

Update the hosts inventory in OMIVV and retry the add host task in the SDDC Manager UI. See the Dell documentation for information about updating the hosts inventory in OMIVV.

The vSAN Performance Service is not enabled for vSAN clusters when CEIP is not enabled

If you do not enable the VMware Customer Experience Improvement Program (CEIP) in SDDC Manager, when you create a workload domain or add a vSphere cluster to a workload domain, the vSAN Performance Service is not enabled for vSAN clusters. When CEIP is enabled, data from the vSAN Performance Service is provided to VMware and this data is used to aid VMware Support with troubleshooting and for products such as VMware Skyline, a proactive cloud monitoring service. See [Customer Experience Improvement Program](#) for more information on the data collected by CEIP.

Enable CEIP in SDDC Manager. See the [VMware Cloud Foundation Documentation](#). After CEIP is enabled, a scheduled task that enables the vSAN Performance Service on existing clusters in workload domains runs every three hours. The service is also enabled for new workload domains and clusters. To enable the vSAN Performance Service immediately, see the [VMware vSphere Documentation](#).

Creation or expansion of a vSAN cluster with more than 32 hosts fails

By default, a vSAN cluster can grow up to 32 hosts. With large cluster support enabled, a vSAN cluster can grow up to a maximum of 64 hosts. However, even with large cluster support enabled, a creation or expansion task can fail on the subtask **Enable vSAN on vSphere Cluster**.

1. Enable Large Cluster Support for the vSAN cluster in the vSphere Client. If it is already enabled skip to step 2.
 - a. Select the vSAN cluster in the vSphere Client.
 - b. Select **Configure > vSAN > Advanced Options**.
 - c. Enable Large Cluster Support.
 - d. Click **Apply**.
 - e. Click **Yes**.
2. Run a vSAN health check to see which hosts require rebooting.
3. Put the hosts into Maintenance Mode and reboot the hosts.

For more information about large cluster support, see <https://kb.vmware.com/kb/2110081>.

Removing a host from a cluster, deleting a cluster from a workload domain, or deleting a workload domain fails if Service VMs (SVMs) are present

If you deployed an endpoint protection service (such as guest introspection) to a cluster through NSX Data Center, then removing a host from the cluster, deleting the cluster, or deleting the workload domain containing the cluster will fail on the subtask **Enter Maintenance Mode on ESXi Hosts**.

- For host removal: Delete the Service VM from the host and retry the operation.
- For cluster deletion: Delete the service deployment for the cluster and retry the operation.
- For workload domain deletion: Delete the service deployment for all clusters in the workload domain and retry the operation.

vCenter Server overwrites the NFS datastore name when adding a cluster to a VI workload domain

If you add an NFS datastore with the same NFS server IP address, but a different NFS datastore name, as an NFS datastore that already exists in the workload domain, then vCenter Server applies the existing datastore name to the new datastore.

If you want to add an NFS datastore with a different datastore name, then it must use a different NFS server IP address.

VMware Cloud Foundation 5.2 on Dell VxRail Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Available Languages](#)
- [Deprecation Notices](#)
- [VMware Cloud Foundation Bill of Materials \(BOM\)](#)
- [Documentation](#)
- [Installation and Upgrade Information](#)
- [Resolved Issues](#)
- [Known Issues](#)

Introduction

VMware Cloud Foundation 5.2 | 15 AUG 2024 | Build 24108943

Check for additions and updates to these release notes.

What's New

The VMware Cloud Foundation (VCF) 5.2 release includes the following:

- **Support for Identity Federation with Entra ID:** VCF users can now configure Microsoft Entra ID (formerly known as Azure AD) as an identity provider.
- **APIs for Auditing PCI Compliance:** VCF users can now use a new set of APIs that audit VCF configuration for compliance with 9 relevant PCI-DSS controls.
- **Avi Load Balancer Integration with VCF:** VCF users can now deploy Avi (formerly NSX Advanced Load Balancer) as part of a new workload domain and perform password rotation and certificate management of the ALB infrastructure from SDDC Manager.
- **Flexible Target BOM for Upgrades:** VCF users can now create a composite and customized BOM using patches when upgrading workload domains. Customers can plan an upgrade along with patches in one orchestrated workflow instead of performing an upgrade and applying patches in separate maintenance windows.
- **Async Patching with SDDC Manager:** Customers previously used the standalone Async Patch Tool to apply patches to the VCF BOM components. VCF 5.2 provides the ability to apply BOM component patches from the SDDC Manager UI.
- **Day N workflows with Embedded Async Patching:** VCF users can now add new workload domains and clusters with patched versions of individual BOM components from SDDC Manager.
- **Asynchronous SDDC Manager Upgrades:** VCF users can now upgrade SDDC Manager independently from the rest of the BOM to apply critical fixes, security patches, and to enable specific features related to SDDC Manager.
- **Authenticated Proxy:** VCF users can now use proxy authentication from SDDC Manager to enable secure connectivity from SDDC Manager to the internet.
- **Offline Depot:** VCF users can now perform lifecycle bundle downloads in offline/air-gapped environments in a simplified manner. The offline depot downloads and stages VCF SDDC Manager and BOM component bundles and enables customers to configure SDDC Manager to download the bundles directly from the offline depot.
- **Isolated Workload Domains Sharing NSX:** VCF users can now create and manage isolated workload domains that can share an NSX Manager instance between them.

Available Languages

Beginning with the next major release, VCF will be supporting the following localization languages:

- Japanese
- Spanish
- French

The following languages will no longer be supported:

- Italian, German, and Simplified Chinese.

Impact:

- Customers who have been using the deprecated languages will no longer receive updates in these languages.
- All user interfaces and help documentation will be available only in English or in the three supported languages mentioned above.

Because VCF localization utilizes the browser language settings, ensure that your settings match the desired language.

Deprecation Notices

- VMware End Of Availability of Perpetual Licensing and SaaS Services. See <https://blogs.vmware.com/cloud-foundation/2024/01/22/vmware-end-of-availability-of-perpetual-licensing-and-saas-services/> for more information.
- In a future release, the "Connect Workload Domains" option from the VMware Aria Operations card located in **SDDC Manager > Administration > Aria Suite** section will be removed and related VCF Public API options will be deprecated.

Starting with VMware Aria Operations 8.10, functionality for connecting VCF Workload Domains to VMware Aria Operations is available directly from the UI. Users are encouraged to use this method within the VMware Aria Operations UI for connecting VCF workload domains, even if the integration was originally set up using SDDC Manager.

- Deprecation announcements for VMware NSX. See the [VMware NSX 4.2.0 Release Notes](#) for details.
 - NSX Manager APIs and NSX Advanced UIs
 - NSX embedded (NSXe)
 - Some parameters in Switch IPFIX
 - NSX Migration Coordinator

VMware Cloud Foundation Bill of Materials (BOM)

The VMware Cloud Foundation software product is comprised of the following software Bill-of-Materials (BOM). The components in the BOM are interoperable and compatible.

Software Component	Version	Date	Build Number
Cloud Builder VM	5.2	23 JUL 2024	24108943
SDDC Manager	5.2	23 JUL 2024	24108943
VxRail Manager	8.0.300	15 AUG 2024	NA
VMware vCenter Server Appliance	8.0 Update 3a	18 JUL 2024	24091160
VMware vSAN Witness Appliance	8.0 Update 3	25 JUN 2024	24022510
VMware NSX	4.2	23 JUL 2024	24105817
VMware Aria Suite Lifecycle	8.18	23 JUL 2024	24029603

- VMware ESXi and VMware vSAN are part of the VxRail BOM.
- You can use VMware Aria Suite Lifecycle to deploy VMware Aria Automation, VMware Aria Operations, VMware Aria Operations for Logs, and Workspace ONE Access (formerly known as VMware Identity Manager). VMware Aria Suite Lifecycle determines which versions of these products are compatible and only allows you to install/upgrade to supported versions. See [VMware Aria Suite Upgrade Paths on VMware Cloud Foundation 4.4.x +](#).
- VMware Aria Operations for Logs content packs are installed when you deploy VMware Aria Operations for Logs.
- The VMware Aria Operations management pack is installed when you deploy VMware Aria Operations.
- You can access the latest versions of the content packs for VMware Aria Operations for Logs from the VMware Solution Exchange and the VMware Aria Operations for Logs in-product marketplace store.

Documentation

The following documentation is available:

- [VMware Cloud Foundation on Dell VxRail Guide](#)
- [VMware Cloud Foundation 5.2 Release Notes](#)
- [Support Matrix of VMware Cloud Foundation on Dell VxRail](#)

Installation and Upgrade Information

You can perform a sequential or skip level upgrade to VMware Cloud Foundation 5.2 on Dell VxRail from VMware Cloud Foundation 4.5 or later. If your environment is at a version earlier than 4.5, you must upgrade the management domain and all VI workload domains to VMware Cloud Foundation 4.5 and then upgrade to VMware Cloud Foundation 5.2.

IMPORTANT:

- You must have VxRail 7.0.410 or later to upgrade to VxRail 8.0.300. If you are on a version earlier than VxRail 7.0.410, you must upgrade to 7.0.410 before upgrading to 8.0.300.
- Before you upgrade a vCenter Server, take a file-based backup. See [Manually Back Up vCenter Server](#).

NOTE: Scripts that rely on SSH being activated on ESXi hosts will not work after upgrading to VMware Cloud Foundation 5.2, since VMware Cloud Foundation 5.2 deactivates the SSH service by default. Update your scripts to account for this new behavior. See [KB 86230](#) for information about activating and deactivating the SSH service on ESXi hosts.

Resolved Issues

The following issues have been resolved:

- The SDDC Manager UI displays the "Use Existing NSX instance" option when sharing is not supported.
- Cannot host an Edge cluster on a vSphere cluster that is created with multiple VDSes where VM MANAGEMENT and MANAGEMENT networks are on different VDSes.
- Adding a VxRail cluster using the SDDC Manager UI displays a validation error.
- Incorrect warning displays for Create Workload Domain and Add Cluster UI.
- When adding a VxRail cluster, the teaming polices are not the same as specified in the request payload.
- Configure DNS/NTP task fails at `Generate Inputs for vRealize Suite Resource Types`

Known Issues

For VMware Cloud Foundation 5.2 known issues, see [VMware Cloud Foundation 5.2 known issues](#). Some of the known issues may be for features that are not available on VMware Cloud Foundation on Dell VxRail.

VMware Cloud Foundation 5.2 on Dell VxRail Known Issues appear below:

SDDC Manager fails on host discovery

VxRail host discovery fails in the SDDC Manager UI with the following error: "Failed to load Host cluster details".
Workaround: Add the VxRail host(s) via the API Explorer. See <https://knowledge.broadcom.com/external/article?articleNumber=379202>.

Error while retrying failed VxRail upgrade to 5.2


Retrying a failed VxRail upgrade fails with the error message `Failed to process the request body`.
Workaround: See <https://www.dell.com/support/kbdoc/000227705>.

Compatibility warning when upgrading from VMware Cloud Foundation 5.1 to 5.2

When you plan an upgrade from VMware Cloud Foundation 5.1 to VMware Cloud Foundation 5.2 a compatibility warning may display.

 VMware Cloud Foundation 5.2.0.0 with VxRail Manager 8.0.300 

Select Version

 Unable to verify the compatibility for the following product versions. Please check the product documentation before proceeding to upgrade: .

Workaround: This upgrade path is fully-supported and you can ignore the warning.

Once VCF has been upgraded to 5.2, using the Bundle Transfer Utility to upload VxRail patches fails

VCF 5.2 and later require a new metadata file (`vxrailPartnerBundleMetadata.json`) which contains information about supported patches. The Bundle Transfer Utility does not upload this file correctly, which prevents future VxRail patching.

Workaround:

1. On the SDDC Manager appliance, replace the `partnerBundleMetadata.json` with `vxrailPartnerBundleMetadata.json`. For example:

```
mv /nfs/vmware/vcf/nfs-mount/bundle/depot/local/vxrailPartnerBundleMetadata.json /
nfs/vmware/vcf/nfsmount/bundle/depot/local/partnerBundleMetadata.json
```

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/bundle/depot/local
```

```
chmod -R 755 /nfs/vmware/vcf/nfs-mount/bundle/depot/local
```

2. Upload the `partnerBundleMetadata` file to LCM using the API. See https://developer.broadcom.com/xapis/vmware-cloud-foundation-for-vxrail-api/latest/bundles/#_usecase_uploadpartnerbundle.

Error during thumbprint verification when adding hosts to a cluster

When you add hosts to a VxRail cluster that is part of a workload domain that uses vLCM images, you may see the following error: Dependency on Enabling SSH could not be performed on one or many host in the request payload.

Workaround: Enable SSH on the host(s) and retry the task.

Error upgrading VxRail Manager to 8.0.300

If the management domain is using vSphere Lifecycle Manager (vLCM) images and a VI workload domain is using vLCM baselines, upgrading to VxRail 8.0.300 may fail with the error `Failure occurred while running an upgrade for bundle: VXRAIL_COMPOSITE-SLIM-8.0.300-XXXXX.Trigger set customised depot meets exception. meet error in vlcm service request exchange`.

Workaround: See <https://www.dell.com/support/kbdoc/0002277289>.

VxRail first run fails with error "Failed to do vLCM remediation for cluster"

If an NVMe device or controllers are not VMware Certified, the VxRail first run fails with a Skyline Health alert.

Workaround: Browse to the Skyline Health section for the cluster using the vSphere Client, silence the alert, and retry the first run.

Creating a VI workload domain or adding a cluster with VMFS on FC storage fails

When you use the SDDC Manager UI to create a VI workload domain or add a cluster with VMFS on FC storage, the task fails.

Workaround: Use the Workflow Optimization script or VMware Cloud Foundation on Dell VxRail API to create the VI workload domain or add the cluster.

VxRail plugin missing after the vCenter Server certificate or password is rotated from SDDC Manager

When you rotate a vCenter Server certificate or password, the VxRail plugin may disappear for clusters managed by another vCenter Server.

Workaround: See <https://www.dell.com/support/kbdoc/000224478>.

Creating a workload domain or adding a cluster using ESXi hosts with GPU drivers

If your ESXi hosts include a GPU driver, you must upload the GPU driver to the VxRail Manager before you:

- Create a VI workload domain that uses vSphere Lifecycle Manager images.

- Add a cluster to a VI workload domain that uses vSphere Lifecycle Manager images.

Workaround: See <https://www.dell.com/support/kbdoc/en-in/000202491>.

VxRail does not allow more than 2 active/one active uplink, 1 standby nics.

Mapping more than 2 uplinks causes VxRail cluster validation to fail during the VxRail Dry Run.

Workaround: When creating a custom NIC profile, map no more than 2 uplinks to active/standby uplinks.

Support for consuming VxRail upgrade bundles for 5x-5y is unavailable

VxRail Manager 8.0.x to 8.0.300 upgrade fails stating a 8.0.x.zip file does not exist error.

Workaround: See KB article [94747](#) for the scripts and manual steps to mitigate this compatibility gap.

Failed VxRail first run prevents new cluster/workload domain creation

If the VxRail first run fails, some objects associated with the failed task remain in the vCenter Server inventory and prevent new cluster/workload domain creation.

Workaround: Remove the inventory objects associated with the failed task using the vSphere Client.

1. Log in to the vSphere Client.
2. In the Hosts and Clusters inventory, right-click the failed cluster and select **vSAN > Shutdown cluster**.
3. After the shutdown completes, right-click the failed cluster and select **Delete**.

After the inventory is cleaned up, you can retry adding a cluster or creating a workload domain.

Add VxRail hosts validation fails

When adding VxRail hosts to a workload domain or cluster that uses Fibre Channel (FC) storage, the task may fail with the message No shared datastore can be found on host. This can happen if you used the Workflow Optimization Script to deploy the workload domain or cluster and chose an FC datastore name other than the default name.

Workaround: Use the VMware Host Client to rename the FC datastore on the new VxRail hosts to match the name you entered when creating the workload domain or cluster. Once the FC datastore name of the new hosts matches the existing FC datastore name, retry the Add VxRail Hosts operation.

Unsupported versions of VxRail not restricted during create cluster/domain operations

VCF on VxRail does not restrict using unsupported/unpaired versions of VxRail for create cluster/domain operations. If nodes are re-imaged with a VxRail version that is not paired with the current VCF release, VCF does not restrict using these nodes for creating a cluster/domain.

Workaround: Use the VxRail Manager version paired with the correct VCF release for create domain/cluster operations.

vSAN/vMotion network disruption can occur when using the workflow optimization script

When you use the workflow optimization script to create a new VI workload domain or add a new cluster to an existing workload domain, you can cause a network disruption on existing vSAN/vMotion networks if:

- The IP range for the new vSAN network overlaps with the IP range for an existing vSAN network.
- The IP range for the new vMotion network overlaps with the IP range for an existing vMotion network.

Workaround: None. Make sure to provide vSAN/vMotion IP ranges that do not overlap with existing vSAN/vMotion networks.

Async Patch Tool Release Notes

This document contains the following sections

- [Introduction](#)
- [What's New](#)
- [Resolved Issues](#)
- [Known Issues](#)

Introduction

VMware Cloud Foundation

Async Patch Tool 1.2 | 23 JUL 2024 | Build 24090705

Check for additions and updates to these release notes.

What's New

The Async Patch Tool is a utility that allows you to apply critical patches to certain VMware Cloud Foundation components (NSX Manager, vCenter Server, and ESXi) outside of VMware Cloud Foundation releases. The Async Patch Tool also supports ESXi and VxRail Manager patching of VMware Cloud Foundation on VxRail and is supported with VMware Cloud Foundation 4.2.1 and later.

IMPORTANT: Starting with VMware Cloud Foundation 5.2, you should apply async patches directly from the SDDC Manager UI. See ["Patching the Management and Workload Domains"](#).

If you are upgrading an async patched system from VMware Cloud Foundation 4.x to 4.y, you must use the Async Patch Tool to enable the upgrade (`-r, --enableVCFUpgrade`). If you are upgrading an async patched system from VMware Cloud Foundation 4.x to 5.x or 5.x to 5.x, you do not need to use the Async Patch Tool to enable the upgrade. For more information, see the [Async Patch Tool documentation](#).

See [KB 88287](#) for information about which async patches are supported with your version of VMware Cloud Foundation. The Knowledge Base article also includes information about supported upgrade paths for VMware Cloud Foundation instances that include an async patch.

Resolved Issues

- Async Patch Tool `-l,--listAsyncPatch` option fails: Using the `--productType, --ptype` option with the `-l, --listAsyncPatch` option fails, unless you also provide a `--sku`.

Known Issues

The stand-alone postcheck option fails for VxRail patches

When you run the standalone postcheck option (`--post, --postcheck`) for a VxRail async patch, the postcheck fails. Workaround: None. However, when you run Async Patch Tool with the enable patch option (`-e, --enableAsyncPatch`) the postchecks are run automatically. If enabling a VxRail async patch succeeds, the postcheck is also considered successful.

Workload Domain enabled 7.0.451 async patch failed with error "Incompatible products found"

On environments with SDDC Manager updated to 5.0, enabling VxRail Manager patches may have issues writing AP tool Interoperability validations (Writing ESXi product).

Workaround:

1. Login to SDDC Manager.
2. Switch to root user.
3. Change the directory to `/opt/vmware/vcf/lcm/lcm-app/conf`.
4. In the LifeCycle Management `application-prod.properties` file add the following property (to exclude ESX_HOST in the supported product types):

```
vcf.compatibility.check.supported.product.types=SDDC_MANAGER,VCENTER,NSX_T_MANAGER,VX_MANAGER
```

5. Restart LCM .

Restart the enable patch workflow from AP tool.

Bundle clean up script fails to clean up async patch bundle on vLCM based clusters

AP Tool `disable all` fails to clean up async patch bundle on vLCM based clusters with error `BUNDLE_CLEANUP_FAILED`.

Workaround: See [KB 89719](#).

Async patch is not available to apply in the SDDC Manager UI

After you use the Async Patch Tool to enable a patch and successfully upload the patch to the internal LCM repository on the SDDC Manager appliance, you may not be able to apply the patch. This can happen if a workload domain is in a failed state.

Workaround: In the SDDC Manager UI, perform a precheck on all the workload domains where you intend to apply the patch and resolve any reported issues. After resolving the issues, the async patch bundle should become available to apply.

The SDDC Manager UI displays an unexpected source version when upgrading SDDC Manager

After enabling upgrade for a VMware Cloud Foundation instance that includes an async patch, the SDDC Manager UI displays an unexpected source version for SDDC Manager. For example, if you apply an async patch to your VMware Cloud Foundation 4.2.1 instance, and then you enable an upgrade to VMware Cloud Foundation 4.4.1.1, the SDDC Manager shows 4.4.1.0 as the source version (instead of 4.2.1.0).



Workaround: None. This is a cosmetic issue and has no impact on the upgrade.

It is not clear in the UI which bundle is the async patch bundle when other SDDC Manager bundles are available

If other SDDC Manager bundles are uploaded onto the SDDC Manager appliance, it might be displayed along with the async patch bundle you have enabled (uploaded) using the Async Patch Tool on the "Available Updates" section of the UI. The async patch bundle might have a similar title to the other bundles. As a result, it might be harder to locate.

Workaround: The uploaded async patch bundle can be identified by the following:

- It will be the only "VMware Software Update" bundle in the list.
- The bundle details have a Bundle ID with a suffix of `-apTool` to signify it is an async patch enabled by Async Patch Tool.

Async Patch Tool fails with `FAILED_VCF_PERMISSIONS_ON_SDDC_OUTPUT_DIRECTORY` or `RUNNING_ROOT_OPERATIONS_FAILED` when running on a security-hardened SDDC Manager

When running certain Async Patch workflows on an SDDC Manager that have been hardened following the VMware Cloud Foundation Security Technical Implementation Guide (STIG), the Async Patch Tool fails when attempting to run certain operations as root user. The AP tool logs mention the tool `Received non-empty output for command that expect empty output`.

Workaround: Contact VMware Support.

Async patch bundles display non-standard version numbers in the SDDC Manager UI

Async patch bundles include non-standard version numbers wherever information about the bundles is displayed in the SDDC Manager UI. For example, `Version 1.1.1-000001` or `Required Version 1.2.0-123456` as seen below.

Additional Bundle Details

Version	1.1.1-000001
Severity	Critical
Vendor	VMware
Bundle ID	d214e445-8509-4d50-adac-59d56acd86ae-apTool
▼ Software Component 1	ESX Server
Description	VMware ESXi Server Update Bundle
Update to Version	7.0.3-19193900
Required Version	1.2.0-123456
Release Date	Feb 10, 2022
Vendor	VMware

Async Patch Bundle Versioning

Version	1.1.1-<xxxxxx>
Product Version 1.1.1.1	1.1.1.1
Required Version: NSX Manager async patch bundle	1.1.0-<xxxxxx>
Required Version: ESXi async patch bundle	1.2.0-<xxxxxx>
Required Version: vCenter Server async patch bundle	1.3.0-<xxxxxx>

NOTE

The non-standard version numbering does not apply to VxRail.

Workaround: None. This is by design and ensures that async patches are prioritized and applied in the correct order.

Update history information for workload domains does not contain all updates

When you deactivate all async patches from the SDDC Manager appliance, any update history for previously enabled or applied async patches is lost. No update history will be visible from the SDDC Manager UI or in the VMware Cloud Foundation API response. Deactivating all patches happens implicitly when you run the Async Patch Tool with the enable VCF upgrade option (`-r, --enableVCFUpgrade`). If you previously enabled an async patch, you must disable all patches before you can run the Async Patch Tool with the enable patch option (`-e, --enableAsyncPatch`) again.

Workaround: View the Async Patch Tool `upgrade_history` logs to review the entire async patch update history. Logs are located in the `/var/log/vmware/vcf/lcm/tools/asyncpatchtool` directory on the SDDC Manager appliance.

Older install or upgrade bundles appear as available to download

If SDDC Manager is connected to the VMware Depot, and you enable an async patch, older bundles, that are not required, may appear as available for download in the SDDC Manager UI (**Lifecycle Management > Bundle Management > Bundles**). For example, if you enable an async patch for vCenter Server 7.0 Update 3d, the bundle for vCenter Server 7.0 Update 2c may appear as available for download.

Workaround: Remove the bundles that you do not require.

1. Get the bundle ID for the bundle you want to remove.
 - a. In the SDDC Manager UI, browse to **Lifecycle Management > Bundle Management > Bundles**.
 - b. Find the bundle you want to remove and click **View Details**.

- c. Copy the bundle ID.
2. SSH in to the SDDC Manager appliance using the `vcf` user account.
3. Enter `su` to switch to the root user.
4. Enter the following command, replacing `<bundle id>` with the bundle ID from step 1: `python /opt/vmware/vcf/lcm/lcm-app/bin/bundle_cleanup.py <bundle id>`

Design Guide

A design model for VMware Cloud Foundation (also called VCF) that is based on industry best practices for SDDC implementation.

The *VMware Cloud Foundation Design Guide* provides the supported design options for VMware Cloud Foundation, and a set of decision points, justifications, implications, and considerations for building each component.

Intended Audience

This *VMware Cloud Foundation Design Guide* is intended for cloud architects who are familiar with and want to use VMware Cloud Foundation to deploy and manage an SDDC that meets the requirements for capacity, scalability, backup and restore, and extensibility for disaster recovery support.

Before You Apply This Guidance

The sequence of the VMware Cloud Foundation documentation follows the stages for implementing and maintaining an SDDC.

To apply this *VMware Cloud Foundation Design Guide*, you must be acquainted with the *Getting Started with VMware Cloud Foundation* documentation and with the *VMware Cloud Foundation Release Notes*. See [VMware Cloud Foundation documentation](#).

For performance best practices for vSphere, see [Performance Best Practices for VMware vSphere 8.0 Update 1](#).

Design Elements

This *VMware Cloud Foundation Design Guide* contains requirements and recommendations for the design of each component of the SDDC. In situations where a configuration choice exists, requirements and recommendations are available for each choice. Implement only those of them that are relevant to your target configuration.

Design Element	Description
Requirement	Required for the operation of VMware Cloud Foundation. Deviations are not permitted.
Recommendation	Recommended as a best practice. Deviations are permitted.

VMware Cloud Foundation Deployment Options in This Design

This design guidance is for the all architecture models of VMware Cloud Foundation. By following the guidance, you can examine the design for these deployment options:

- Single VMware Cloud Foundation instance.
- Single VMware Cloud Foundation instance with multiple availability zones (also known as stretched deployment). The default vSphere cluster of the workload domain is stretched between two availability zones by using [vSAN](#) and configuring [vSphere DRS rules](#) and [BGP routing](#) accordingly.
- Multiple VMware Cloud Foundation instances. You deploy several instances of VMware Cloud Foundation to address requirements for scale and co-location of users and resources.

For disaster recovery, workload mobility, or propagation of common configuration to multiple VMware Cloud Foundation instances, you can deploy [NSX Federation](#) for the SDDC management and workload components.

- Multiple VMware Cloud Foundation instances with multiple availability zones. You apply the configuration for stretched clusters for a single VMware Cloud Foundation instance to one or more additional VMware Cloud Foundation instances in your environment.

vCenter Single Sign-On Options in This Design

This design guidance covers the topology with a single vCenter Single Sign-On domain in a VMware Cloud Foundation instance and the topology with several isolated vCenter Single Sign-On domains in a single instance. See [vCenter Single Sign-On Design Requirements for](#) .

VMware Cloud Foundation Design Blueprints

You can follow design blueprints for selected architecture models and topologies that list the applicable design elements. See [VMware Cloud Foundation Topology Design Blueprints](#).

Information about Environment Configurations that Can be Converted or Imported into VMware Cloud Foundation

In VMware Cloud Foundation 5.2, by using the VCF Import Tool, you can convert or import infrastructure into your VMware Cloud Foundation in the following way:

- If you do not already have SDDC Manager deployed, you can deploy it on an existing vSphere environment and use the VCF Import Tool to convert that environment to the VMware Cloud Foundation management domain.
- If SDDC Manager is already deployed, you can use the VCF Import Tool to import existing vSphere environments as VI workload domains.

For more information, see [Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#) in the *VMware Cloud Foundation Administration Guide*.

VMware Cloud Foundation Glossary

See [VMware Cloud Foundation Glossary](#) for constructs, operations, and other terms specific to VMware Cloud Foundation. It is important to understand these constructs before continuing with this design guidance.

VMware Cloud Foundation Concepts

To design a VMware Cloud Foundation deployment, you need to understand certain *VMware Cloud Foundation* concepts.

Architecture Models and Workload Domain Types in VMware Cloud Foundation

When you design a VMware Cloud Foundation deployment, you decide what architecture model, that is, standard or consolidated, and what workload domain types, for example, consolidated, isolated, or standard, to implement according to the requirements for hardware, expected number of workloads and workload domains, co-location of management and customer workloads, identity isolation, and shared or isolated networking and security.

Architecture Models

Decide on a model according to your organization's requirements and your environment's resource capabilities. Implement a standard architecture for workload provisioning and mobility across VMware Cloud Foundation instances according to production best practices. If you plan to deploy a small-scale environment, or if you are working on an SDDC proof-of-concept, implement a consolidated architecture.

Figure 1: Choosing a VMware Cloud Foundation Architecture Model

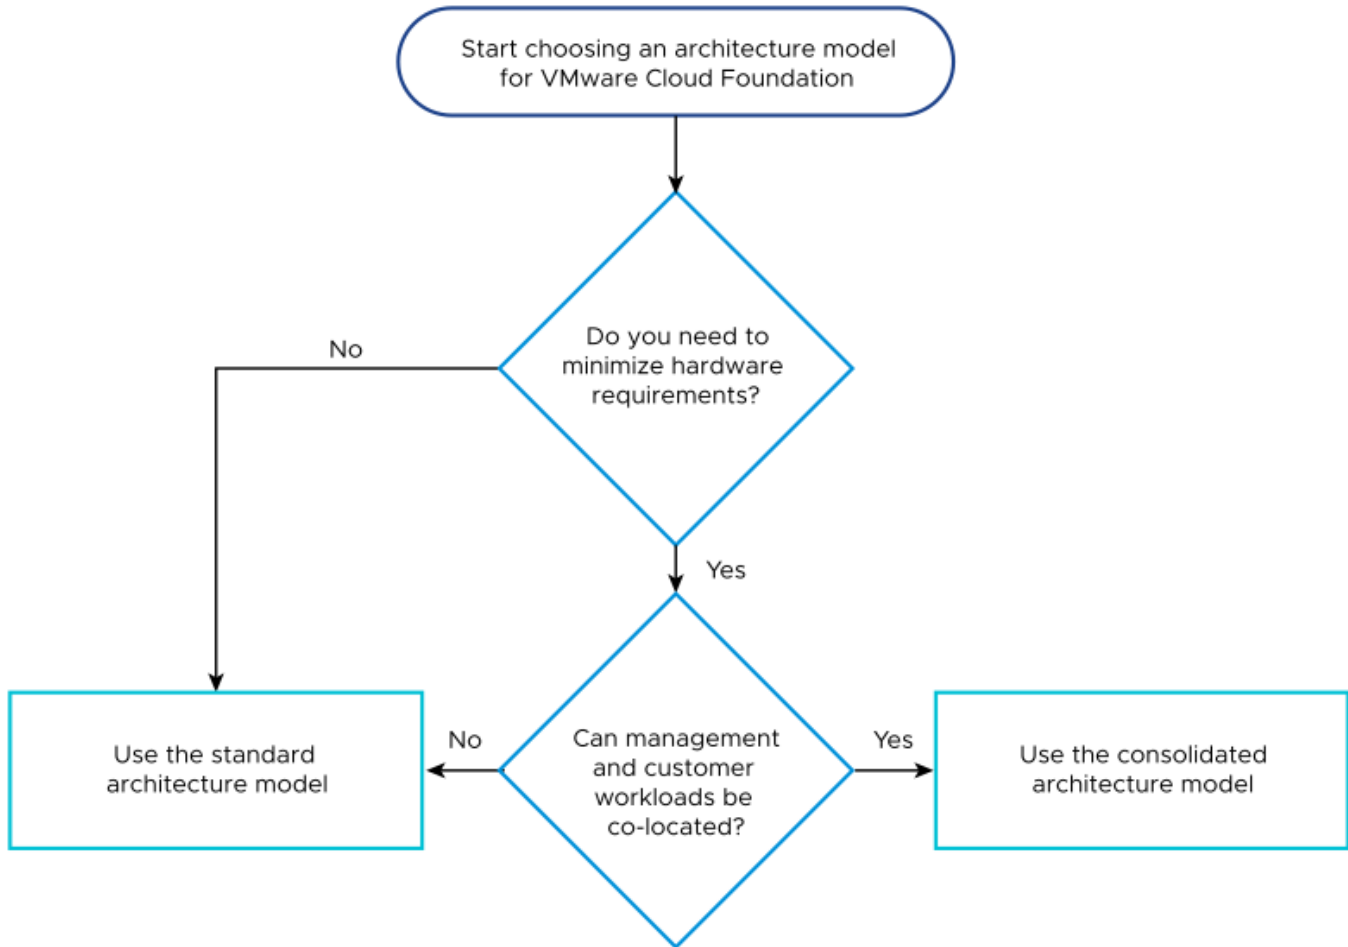


Table 1: Architecture Model Recommendations for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-ARCH-RCMD-CFG-001	Use the standard architecture model of VMware Cloud Foundation.	<ul style="list-style-type: none"> Aligns with the VMware best practice of separating management workloads from customer workloads. Provides better long-term flexibility and expansion options. 	Requires additional hardware.

Workload Domain Types

A workload domain represents a logical unit of application-ready infrastructure that groups ESXi hosts managed by a vCenter Server instance with specific characteristics according to VMware recommended practices. A workload domain can consist of one or more vSphere clusters, provisioned by SDDC Manager.

Table 2: Workload Domain Types

Workload Domain Type	Description	Benefits	Drawbacks
Management domain	<ul style="list-style-type: none"> • First domain deployed. • Contains the following management appliances for all workload domains: <ul style="list-style-type: none"> – vCenter Server – NSX Manager – SDDC Manager – Optional. VMware Aria Suite components – Optional. Management domain NSX Edge nodes • Has dedicated ESXi hosts • First domain to upgrade. 	<ul style="list-style-type: none"> • Guaranteed sufficient resources for management components • Makes it possible to use specific hardware to meet the needs only of the management components • Makes it possible to use dedicated physical compute, network and storage separately from those used for additional workloads • Enables separate lifecycle management of management and workload components. 	<ul style="list-style-type: none"> • You must carefully size the domain to accommodate planned deployment of VI workload domains and additional management components. • Hardware might not be fully utilized until full-scale deployment has been reached.
Consolidated domain	<ul style="list-style-type: none"> • Represents a management domain which also runs customer workloads. • Uses resource pools to ensure sufficient resources for management components. 	<ul style="list-style-type: none"> • Considers the minimum possible initial hardware and management component footprint. • Can be scaled to a standard architecture model. 	<ul style="list-style-type: none"> • Management components and customer workloads are not isolated. • You must constantly monitor it to ensure sufficient resources for management components. • Migrating customer workloads to dedicated VI workload domains is more complex.
VI workload domain	<ul style="list-style-type: none"> • Represents an additional workload domain for running customer workloads. • Shares a vCenter Single Sign-On domain with the management domain. • Shares identity provider configuration with the management domain. • Has dedicated ESXi hosts. 	<ul style="list-style-type: none"> • Can share an NSX Manager instance with other VI workload domains. • All workload domains can be managed through a single pane of glass. • Reduces password management overhead. • Enables independent life cycle management. 	This workload domain type cannot provide distinct vCenter Single Sign-On domains for customer workloads.
Isolated VI workload domain	<ul style="list-style-type: none"> • Represents an additional workload domain for running customer workloads. • Has a distinct vCenter Single Sign-On domain. 	<ul style="list-style-type: none"> • Can provide distinct vCenter Single Sign-On domains for customer workloads. • You can scale up to 24 VI workload domains per 	<ul style="list-style-type: none"> • Workload domain vCenter Server instances are managed through different panes of glass. • Additional password management overhead

Table continued on next page

Continued from previous page

Workload Domain Type	Description	Benefits	Drawbacks
	<ul style="list-style-type: none">• Has a distinct identity provider configuration.• Has dedicated ESXi hosts.	VMware Cloud Foundation instance. <ul style="list-style-type: none">• Enables independent life cycle management.	exists for administrators of VMware Cloud Foundation.

Figure 2: Choosing a VMware Cloud Foundation Workload Domain Type for Customer Workloads

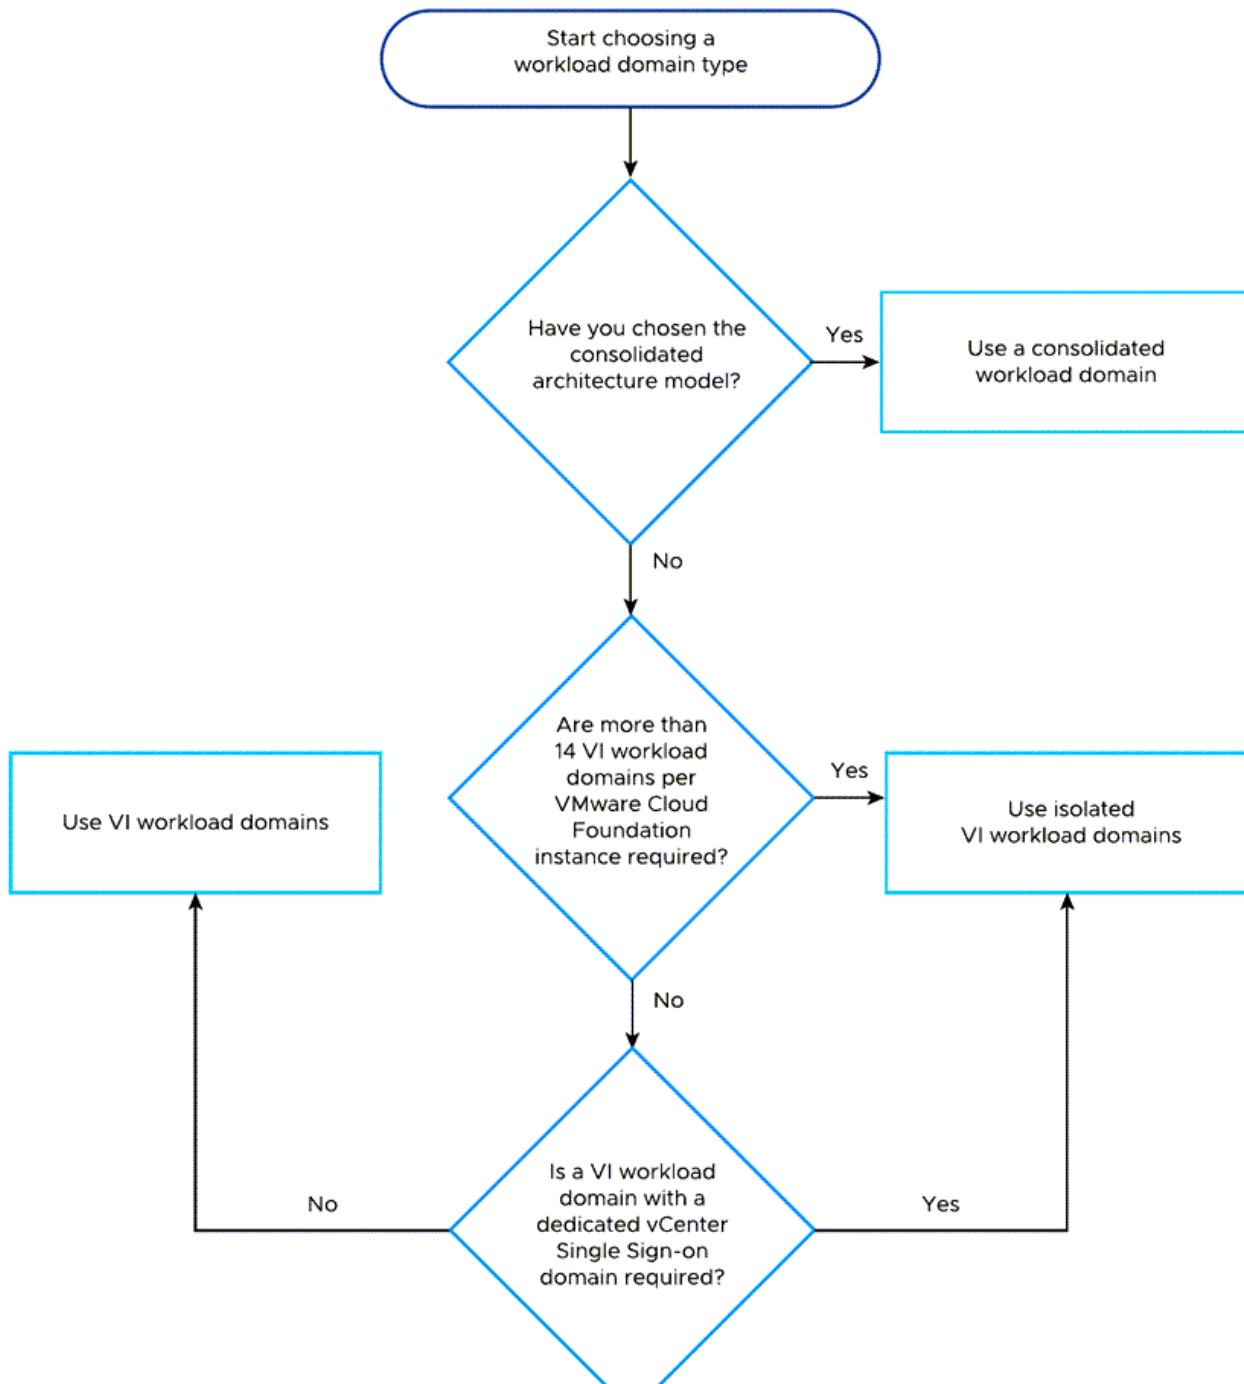


Table 3: Workload Domain Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WLD-RCMD-CFG-001	Use VI workload domains or isolated VI workload	<ul style="list-style-type: none"> Aligns with the VMware best practice of 	Requires additional hardware.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
	domains for customer workloads.	separating management workloads from customer workloads. <ul style="list-style-type: none"> Provides better long term flexibility and expansion options. 	

Workload Domain Cluster to Rack Mapping in VMware Cloud Foundation

VMware Cloud Foundation distributes the functionality of the SDDC across multiple workload domains and vSphere clusters. A workload domain, whether it is the management workload domain or a VI workload domain, is a logical abstraction of compute, storage, and network capacity, and consists of one or more vSphere clusters. Each cluster can exist vertically in a single rack or be spanned horizontally across multiple racks.

The relationship between workload domain clusters and data center racks in VMware Cloud Foundation is not one-to-one. While a workload domain cluster is an atomic unit of repeatable building blocks, a rack is a unit of size. Because workload domain clusters can have different sizes, you map workload domain clusters to data center racks according to your requirements and physical infrastructure constraints. You determine the total number of racks for each cluster type according to your scalability needs.

Table 4: Workload Domain Cluster to Rack Configuration Options

Workload Domain Cluster to Rack Configuration	Description
Workload domain cluster in a single rack	<ul style="list-style-type: none"> The workload domain cluster occupies a single rack. Can be used for shared edge and compute workloads in the same cluster. Can be dedicated for compute-only workloads or for NSX Edge-only clusters.
Workload domain cluster spanning multiple racks	<ul style="list-style-type: none"> The management domain default cluster can span multiple racks if the data center fabric can provide Layer 2 adjacency, such as VXLAN overlay in the fabric, between racks. If the Layer 3 fabric does not support this requirement, then the management default cluster must be mapped to a single rack. A VI workload domain cluster dedicated to compute-only workloads, without SDDC Manager deployed NSX Edge clusters, can span racks when using NSX Overlay in conjunction with a Layer 3 network fabric without Layer 2 adjacency between racks. A vSphere cluster that is to host only an NSX Edge cluster, deployed from SDDC Manager, must be deployed in a single rack. To increase redundancy, you must deploy two vSphere clusters to two separate racks with edge nodes in both racks.
Workload domain cluster with multiple availability zones, each zone in a single rack	<ul style="list-style-type: none"> To span multiple availability zones, the network fabric must support stretched Layer 2 networks and Layer 3 routed networks between the availability zones.

Table continued on next page

Continued from previous page

Workload Domain Cluster to Rack Configuration	Description
	<ul style="list-style-type: none">A cluster spanning multiple racks is not supported with multiple availability zones.

Figure 3: Workload Domain Cluster in a Single Rack




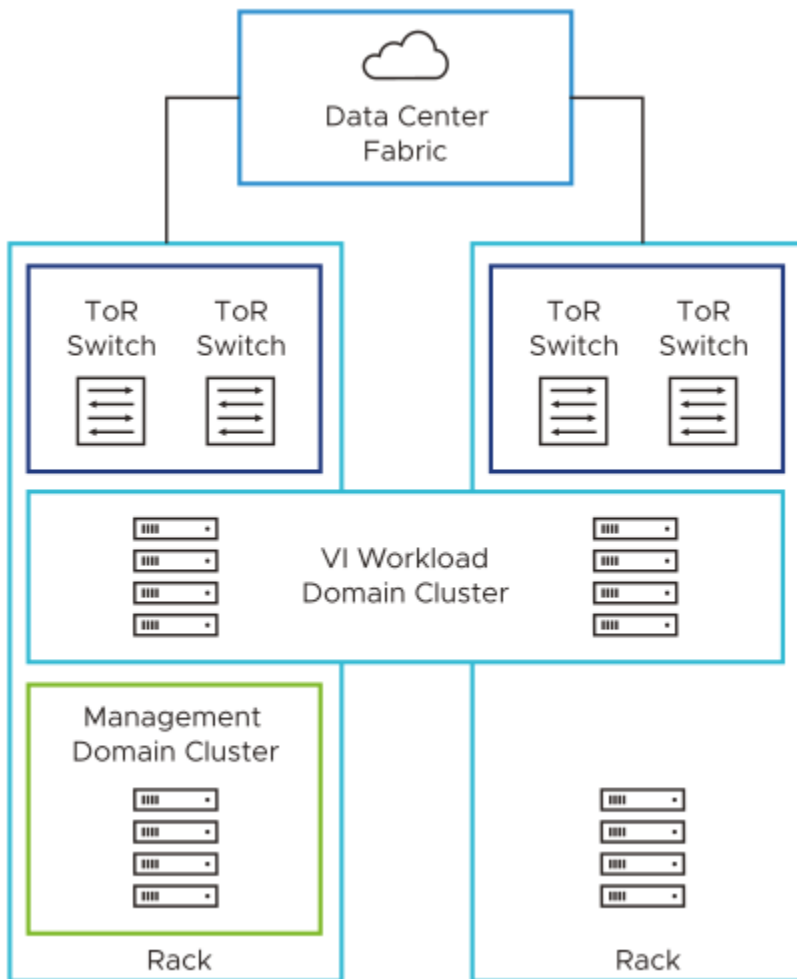
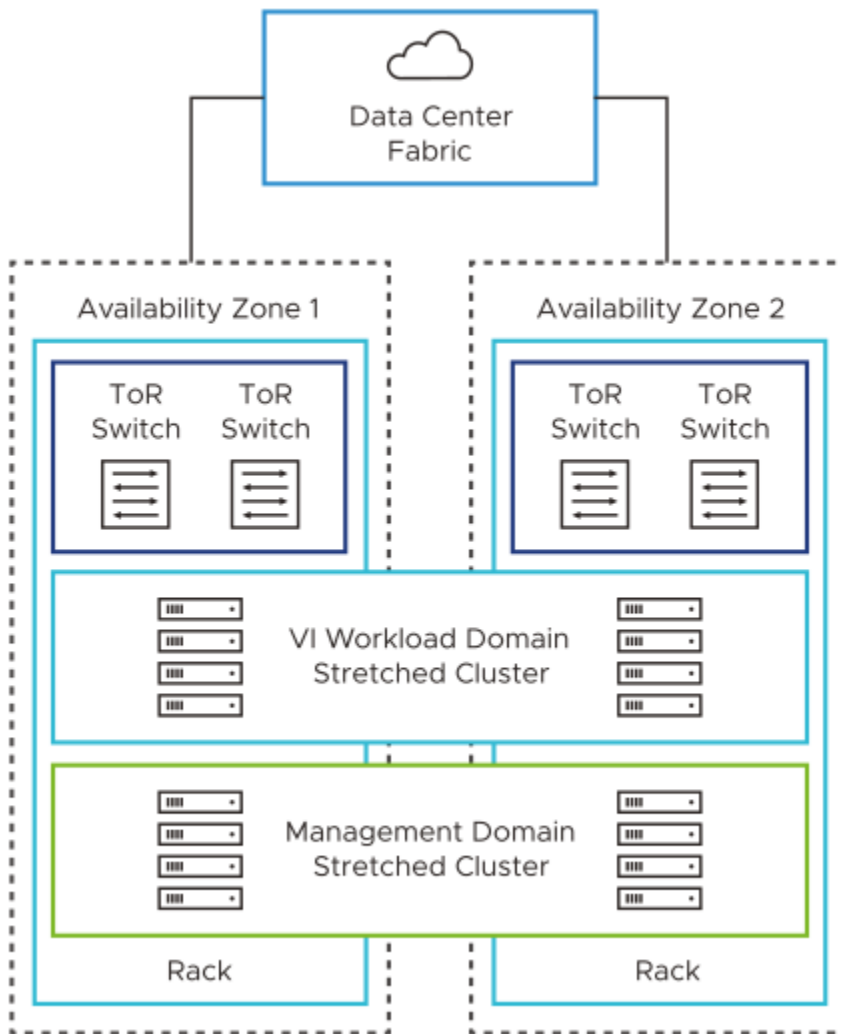
Figure 4: Workload Domain Cluster Spanning Multiple Racks

Figure 5: Workload Domain Cluster with Multiple Availability Zones, Each Zone in One Rack



Networking Models in VMware Cloud Foundation

VMware Cloud Foundation supports multiple networking models that provide different levels of network availability, resilience, and scale.

Networking Model	Description	Benefits	Drawbacks
NSX overlay	Overlay-backed NSX segments	<ul style="list-style-type: none"> Supports AVNs in the management domain. Supports deployment of NSX Edge clusters from SDDC Manager. 	<ul style="list-style-type: none"> Requires routing configured on the physical network fabric to route traffic from the NSX Edge nodes.
NSX VLAN-backed	VLAN-backed NSX segments	<ul style="list-style-type: none"> Deploying NSX Edge clusters is not required. Routing configuration on the ToR switches for NSX 	<ul style="list-style-type: none"> Requires changes in the physical network fabric to add new networks and VLANs.

Table continued on next page

Continued from previous page

Networking Model	Description	Benefits	Drawbacks
		Edge nodes is not required.	

External Services Design for VMware Cloud Foundation

IP addressing scheme, name resolution, and time synchronization must support the requirements for VMware Cloud Foundation deployments.

Table 5: External Services Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-EXT-REQD-NET-001	Allocate statically assigned IP addresses and host names for all workload domain components.	Ensures stability across the VMware Cloud Foundation instance, and makes it simpler to maintain, track, and implement a DNS configuration.	You must provide precise IP address management.
VCF-EXT-REQD-NET-002	Configure forward and reverse DNS records for all workload domain components.	Ensures that all components are accessible by using a fully qualified domain name instead of by using IP addresses only. It is easier to remember and connect to components across the VMware Cloud Foundation instance.	You must provide DNS records for each component.
VCF-EXT-REQD-NET-003	Configure time synchronization by using an internal NTP time source for all workload domain components.	Ensures that all components are synchronized with a valid time source.	An operational NTP service must be available in the environment.
VCF-EXT-REQD-NET-004	Set the NTP service for all workload domain components to start automatically.	Ensures that the NTP service remains synchronized after you restart a component.	None.

Supported Storage Types for VMware Cloud Foundation

Storage design for VMware Cloud Foundation includes the design for principal and supplemental storage.

Principal storage is used during the creation of a workload domain and is capable of running workloads. Supplemental storage can be added after the creation of a workload domain and can be capable of running workloads or be used for data at rest storage such as virtual machine templates, backup data, and ISO images.

Special considerations apply if you plan to add clusters to the management domain, for example, to separate additional management components that require specific hardware resources or might impact the performance of the main management components in the default cluster, or, in the case of the consolidated architecture of VMware Cloud Foundation, to separate customer workloads from the management components.

VMware Cloud Foundation supports the following principal and supplemental storage combinations for clean deployments of VMware Cloud Foundation.

NOTE

For information about the supported storage types for vSphere environments converted or imported into VMware Cloud Foundation, see [Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#) in the *VMware Cloud Foundation Administration Guide*.

Table 6: Supported Storage Types in VMware Cloud Foundation

Storage Type	Management Domain - Default Cluster	Management Domain - Additional Clusters	VI Workload Domain
vSAN Original Storage Architecture (OSA)	Principal	Principal	Principal
VMware vSAN Max™	Not supported	Principal	Principal
Cross-cluster capacity sharing (HCI Mesh)	Supplemental	<ul style="list-style-type: none"> Principal (additional compute-only clusters) Supplemental 	<ul style="list-style-type: none"> Principal (additional compute-only clusters) Supplemental
VMware vSphere® Virtual Volumes™ (FC, iSCSI, or NFS)	Supplemental	Supplemental	<ul style="list-style-type: none"> Principal Supplemental
VMFS on FC	Supplemental	<ul style="list-style-type: none"> Principal (SDDC API only) Supplemental 	<ul style="list-style-type: none"> Principal Supplemental
NFS	Supplemental (NFS 3 and NFS 4.1)	<ul style="list-style-type: none"> Principal (NFS 3 and SDDC API only) Supplemental (NFS 3 and NFS 4.1) 	<ul style="list-style-type: none"> Principal (NFS 3) Supplemental (NFS 3 and NFS 4.1)
iSCSI	Supplemental	Supplemental	Supplemental
NVMe/TCP	Supplemental	Supplemental	Supplemental
NVMe/FC	Supplemental	Supplemental	Supplemental
NVMe/RDMA	Supplemental	Supplemental	Supplemental

NOTE

For a consolidated VMware Cloud Foundation architecture model, the storage types that are supported for the management domain apply.

vSphere Design for VMware Cloud Foundation

The vSphere design includes determining the configuration of the vCenter Server instances, ESXi hosts, vSphere clusters, and vSphere networking for a VMware Cloud Foundation instance.

NSX Design for VMware Cloud Foundation

In VMware Cloud Foundation, you use NSX to implement virtualization for networks, routing and load balancing. It provides support for an automated approach to the creation of virtual network segments and routing objects used to connect management and customer virtual machines to the physical network.

You also create constructs for solutions that are deployed for a single VMware Cloud Foundation instance or are available across multiple VMware Cloud Foundation instances. These constructs provide routing to the data center and load balancing.

Table 7: NSX Logical Concepts and Components

Component	Description
NSX Manager	<ul style="list-style-type: none"> • Provides the user interface and the REST API for creating, configuring, and monitoring NSX components, such as segments, and Tier-0 and Tier-1 gateways. • In a deployment with NSX Federation, NSX Manager is called NSX Local Manager.
NSX Edge nodes	<ul style="list-style-type: none"> • A special type of transport node that contains service router components. • Provides north-south traffic connectivity between the physical data center networks and the NSX SDN networks. Each NSX Edge node has multiple interfaces where traffic flows. • Provides east-west traffic flow between virtualized workloads. Provides stateful services such as load balancing and DHCP. In a deployment with multiple VMware Cloud Foundation instances, east-west traffic between the VMware Cloud Foundation instances flows through the NSX Edge nodes.
NSX Federation (optional design extension)	<ul style="list-style-type: none"> • Propagates configurations that span multiple NSX instances in a single VMware Cloud Foundation instance or across multiple VMware Cloud Foundation instances. You can stretch overlay segments, activate failover of segment ingress and egress traffic between VMware Cloud Foundation instances, and implement a unified firewall configuration. • In a deployment with multiple VMware Cloud Foundation instances, you use NSX to provide cross-instance services to SDDC management components that do not have native support for availability at several locations, such as VMware Aria Automation and VMware Aria Operations. • Connect only workload domains of matching types (management domain to management domain or VI workload domain to VI workload domain).
NSX Global Manager (Federation only)	<ul style="list-style-type: none"> • Is part of deployments with multiple VMware Cloud Foundation instances where NSX Federation is required. NSX Global Manager can connect multiple NSX Local Manager instances under a single global management plane. • Provides the user interface and the REST API for creating, configuring, and monitoring NSX global objects, such as global virtual network segments, and global Tier-0 and Tier-1 gateways. • Connected NSX Local Manager instances create the global objects on the underlying software-defined network that you define from NSX Global Manager. An NSX Local Manager instance directly communicates with other NSX Local Manager instances to synchronize configuration and state needed to implement a global policy. • NSX Global Manager is a deployment-time role that you assign to an NSX Manager appliance.
NSX Manager instance shared between VI workload domains	<ul style="list-style-type: none"> • An NSX Manager instance can be shared between up to 14 VI workload domains that are part of the same vCenter Single Sign-On domain

Table continued on next page

Continued from previous page

Component	Description
	<ul style="list-style-type: none"> • An NSX Manager instance can be shared between up to 16 isolated VI workload domains. • Using a shared NSX Manager instance reduces resource requirements for the management domain. • A single transport zone is shared across all clusters in all VI workload domains that share the NSX Manager instance. • The management domain NSX Manager instance cannot be shared.

Logical Design for NSX for VMware Cloud Foundation

NSX provides networking services to workloads in VMware Cloud Foundation such as load balancing, routing and virtual networking.

Table 8: NSX Logical Design

Component	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
NSX Manager Cluster	<ul style="list-style-type: none"> • Three appropriately sized nodes with a virtual IP (VIP) address with an anti-affinity rule to keep them on different hosts. • vSphere HA protects the cluster nodes applying high restart priority 	<ul style="list-style-type: none"> • Three appropriately sized nodes with a VIP address with an anti-affinity rule to keep them on different hosts. • vSphere HA protects the cluster nodes by applying high restart priority • vSphere DRS rule should-run-on-hosts-in-group keeps the NSX Manager VMs in the first availability zone.
NSX Global Manager Cluster (Conditional)	<ul style="list-style-type: none"> • Manually deployed three appropriately sized nodes with a VIP address with an anti-affinity rule to run them on different hosts. • One active and one standby cluster. • vSphere HA protects the cluster nodes applying high restart priority. 	<ul style="list-style-type: none"> • Manually deployed three appropriately sized nodes with a VIP address with an anti-affinity rule to run them on different hosts. • One active and one standby cluster. • vSphere HA protects the cluster nodes by applying high restart priority. • vSphere DRS rule should-run-on-hosts-in-group keeps the NSX Global Manager VMs in the first availability zone.
NSX Edge Cluster	<ul style="list-style-type: none"> • Two appropriately sized NSX Edge nodes with an anti-affinity rule to separate them on different hosts. • vSphere HA protects the cluster nodes applying high restart priority. 	<ul style="list-style-type: none"> • Two appropriately sized NSX Edge nodes in the first availability zone with an anti-affinity rule to separate them on different hosts. • vSphere HA protects the cluster nodes by applying high restart priority.

Table continued on next page

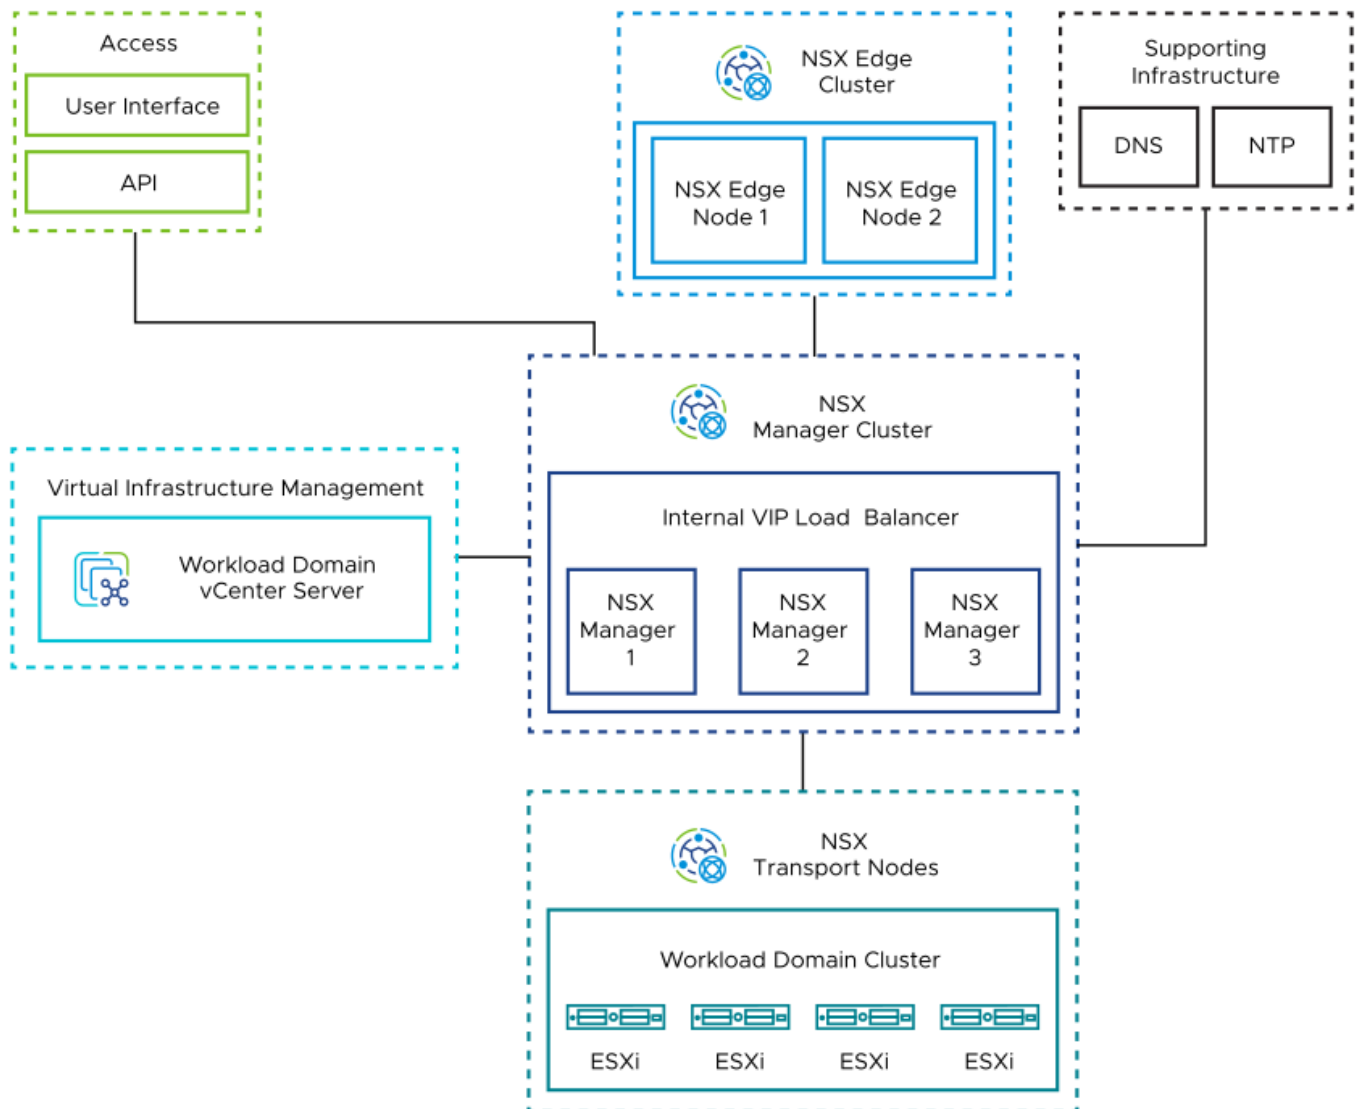
Continued from previous page

Component	VMware Cloud Foundation Instances with a Single Availability Zone	VMware Cloud Foundation Instances with Multiple Availability Zones
		<ul style="list-style-type: none"> vSphere DRS rule should-run-on-hosts-in-group keeps the NSX Edge VMs in the first availability zone.
Transport Nodes	<ul style="list-style-type: none"> Each ESXi host acts as a host transport node. Minimum two edge transport nodes. 	<ul style="list-style-type: none"> Each ESXi host acts as a host transport node. Minimum two transport nodes in the first availability zone.
Transport zones	<ul style="list-style-type: none"> One VLAN transport zone for north-south traffic. Maximum one overlay transport zone for overlay segments per NSX instance. One VLAN transport zone for VLAN-backed segments. 	<ul style="list-style-type: none"> One VLAN transport zone for north-south traffic. Maximum one overlay transport zone for overlay segments per NSX instance. One or more VLAN transport zones for VLAN-backed segments.
VLANs and IP subnets allocated to NSX For information about the networks for virtual infrastructure management, see Distributed Port Group Design .	See VLANs and Subnets for .	See VLANs and Subnets for .
Routing configuration	<ul style="list-style-type: none"> BGP for a single VMware Cloud Foundation instance. In a VMware Cloud Foundation deployment with NSX Federation, BGP with ingress and egress traffic to the first VMware Cloud Foundation instance during normal operating conditions. 	<ul style="list-style-type: none"> BGP with path prepend to control ingress traffic and local preference to control egress traffic through the first availability zone during normal operating conditions. In a VMware Cloud Foundation deployment with NSX Federation, BGP with ingress and egress traffic to the first instance during normal operating conditions.

For a description of the NSX logical component in this design, see [Table : NSX Logical Concepts and Components](#).

Single Instance - Single Availability Zone

The NSX design for the Single Instance - Single Availability Zone topology consists of the following components:

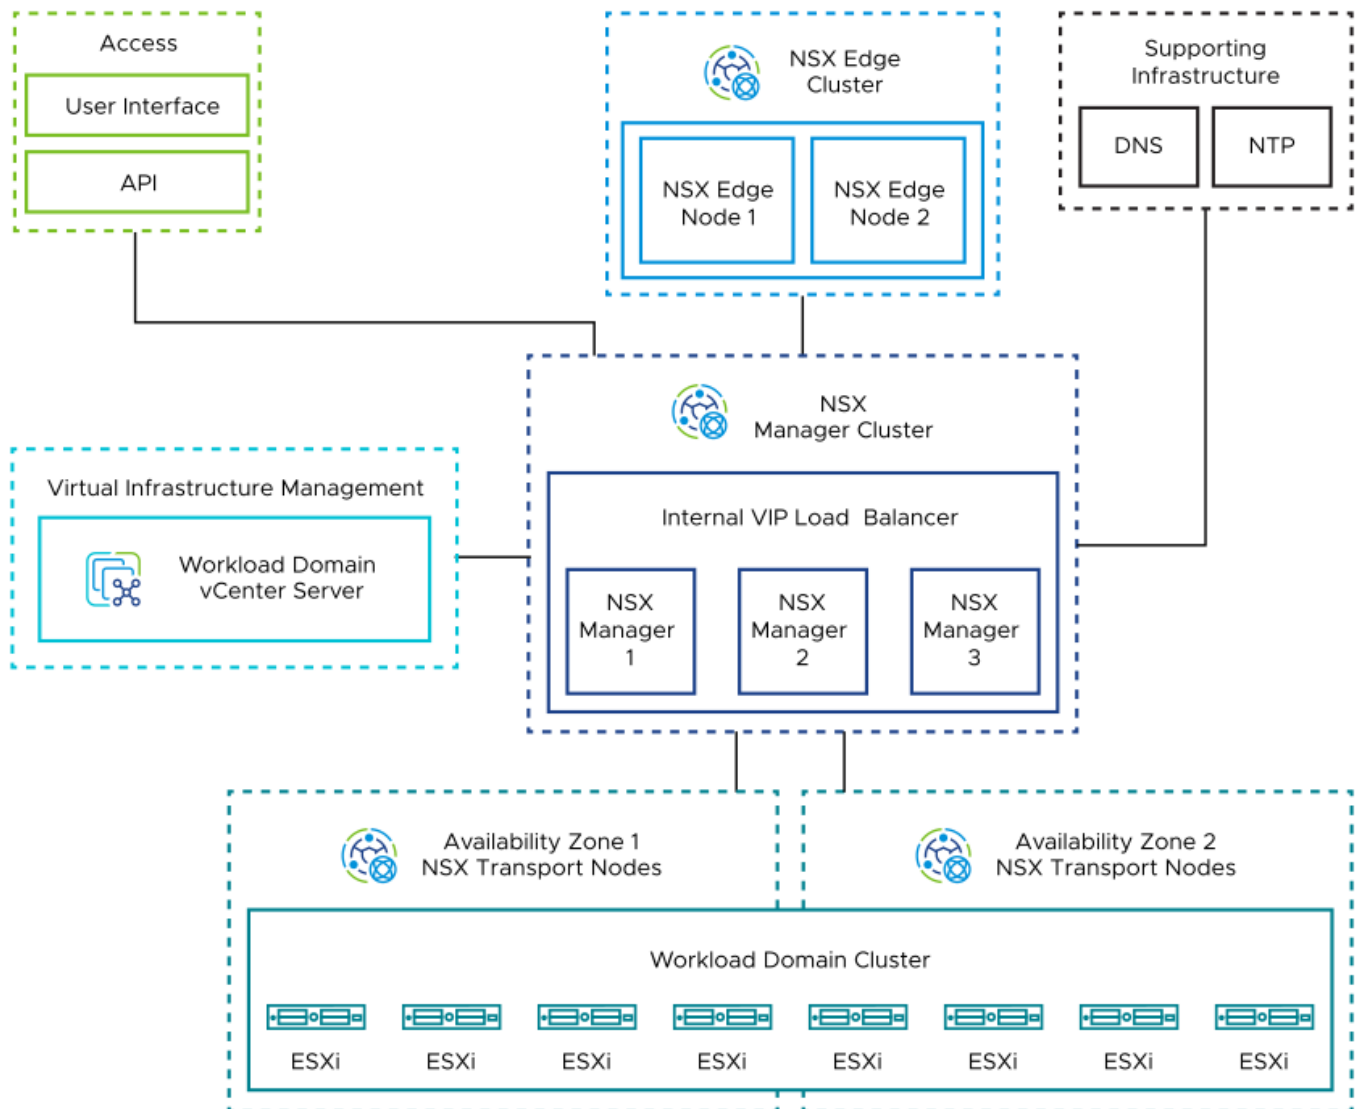
Figure 6: NSX Logical Design for a Single Instance - Single Availability Zone Topology

- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes in the workload domain that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts in the workload domain that are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.

Single Instance - Multiple Availability Zones

The NSX design for a Single Instance - Multiple Availability Zone topology consists of the following components:

Figure 7: NSX Logical Design for a Single Instance - Multiple Availability Zone Topology

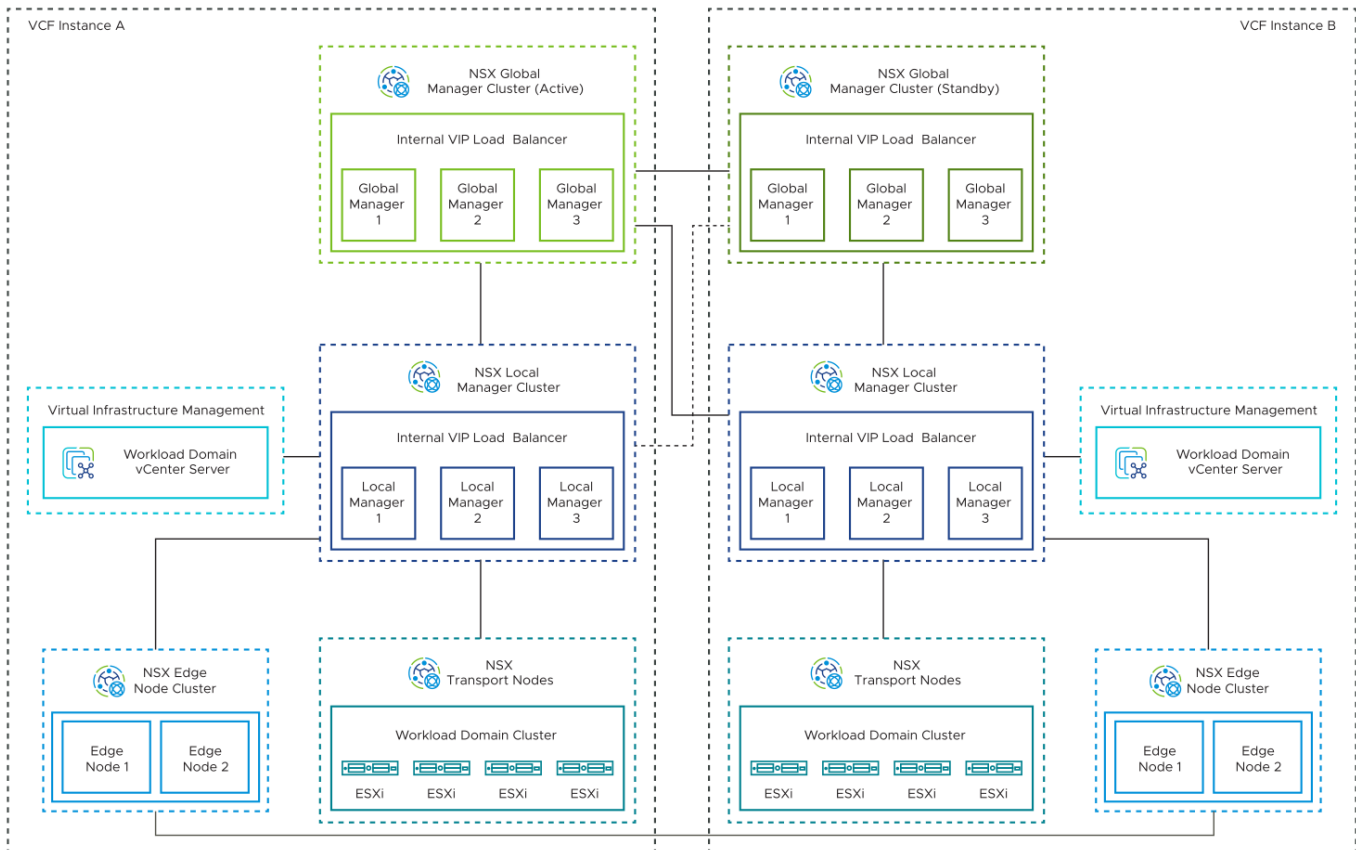


- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts that are distributed evenly across availability zones in the workload domain and are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.

Multiple Instances - Single Availability Zone

The NSX design for a Multiple Instance - Single Availability Zone topology consists of the following components:

Figure 8: NSX Logical Design for a Multiple Instance - Single Availability Zone Topology

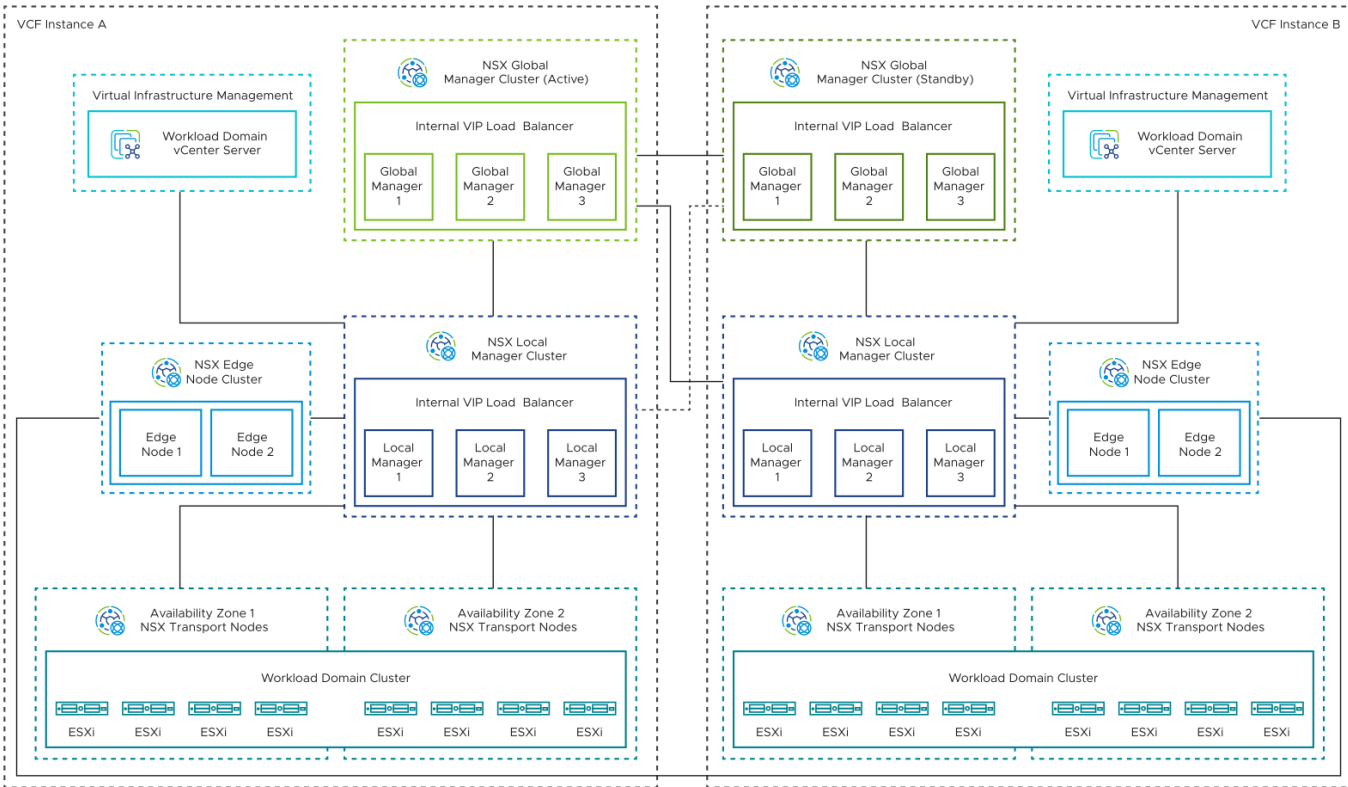


- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts in the workload domain that are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.
- NSX Global Manager clusters deployed in each of the first two VMware Cloud Foundation instances in an Active/Standby configuration. You deploy the NSX Global Manager clusters so that you can use NSX Federation for global management of networking and security services.
- An additional infrastructure VLAN in each VMware Cloud Foundation instance to carry instance-to-instance traffic (RTEP).

Multiple Instances - Multiple Availability Zones

The NSX design for a Multiple Instance - Multiple Availability Zone topology consists of the following components:

Figure 9: NSX Logical Design for Multiple Instance - Multiple Availability Zone Topology



- Unified appliances that have both the NSX Local Manager and NSX Controller roles. They provide management and control plane capabilities.
- NSX Edge nodes that provide advanced services such as load balancing, and north-south connectivity.
- ESXi hosts that are distributed evenly across availability zones in the workload domain in a VMware Cloud Foundation instance, and are registered as NSX transport nodes to provide distributed routing and firewall services to workloads.
- NSX Global Manager cluster in each of the first two VMware Cloud Foundation instances.

You deploy the NSX Global Manager cluster in each VMware Cloud Foundation instance so that you can use NSX Federation for global management of networking and security services.

- An additional infrastructure VLAN in each VMware Cloud Foundation instance to carry instance-to-instance traffic (RTEP).

Routing Design for VMware Cloud Foundation

NSX Edge clusters in VMware Cloud Foundation provide pools of capacity for service router functions in NSX.

Routing Options for VMware Cloud Foundation

VMware Cloud Foundation supports the following routing options:

Routing Type	Description	Benefits	Drawbacks
Static routing	<ul style="list-style-type: none"> • The network administrator manages the routing information, 	<ul style="list-style-type: none"> • No dynamic routing protocol is required on the ToR switches. 	<ul style="list-style-type: none"> • You must manually create static routes in

Table continued on next page

Continued from previous page

Routing Type	Description	Benefits	Drawbacks
	<ul style="list-style-type: none"> adding routing information to the routing table. If any change occurs in the network, the administrator has to update the related information in the routing table. 	<ul style="list-style-type: none"> Might reduce ToR switch license costs. 	<ul style="list-style-type: none"> NSX Manager on the Tier-0 gateway. If required, you must manually create an HA VIP in the NSX Manager on the Tier-0 gateway to provide redundancy across ToR switches. Not supported with vSAN stretched clusters.
OSPF	<ul style="list-style-type: none"> The routing protocol automatically adds and manages the routing information in the routing table. If any change occurs in the network, the routing protocol automatically updates the related information in the routing table. If any new segments or subnets are added in NSX, they are automatically added to the routing table. 	<ul style="list-style-type: none"> If the physical fabric is running OSPF routing protocol, using OSPF at the virtual layer might be a simpler approach for the network administrator. 	<ul style="list-style-type: none"> Needs additional manual configuration. See VMware Knowledge Base article 85916. Not supported with vSAN stretched clusters. Not supported with NSX Federation. Combined use of BGP and OSPF on a single Tier-0 gateway not supported.
BGP	<ul style="list-style-type: none"> BGP is known as an exterior gateway protocol. It is designed to share routing information between disparate networks, known as autonomous systems (ASes). When multiple BGP-derived paths exist, the protocol chooses a path to send traffic based on certain criteria. The routing protocol automatically adds and manages the routing information in the routing table. If any new segments or subnets are added in NSX, they are automatically added to the routing table. 	<ul style="list-style-type: none"> Fully supported by the automated edge workflows in VMware Cloud Foundation. Fully supported for all VMware Cloud Foundation topologies. 	<ul style="list-style-type: none"> None.

BGP routing is the routing option recommended for VMware Cloud Foundation.

BGP Routing Design for VMware Cloud Foundation

Determine the number, networking, and high-availability configuration of the Tier-0 and Tier-1 gateways in NSX for VMware Cloud Foundation workload domains. Identify the BGP configuration for a single availability zone and two availability zones in the environment.

Table 9: Routing Direction Definitions

Routing Direction	Description
North-south	Traffic leaving or entering the NSX domain, for example, a virtual machine on an overlay network communicating with an end-user device on the corporate network.
East-west	Traffic that remains in the NSX domain, for example, two virtual machines on the same or different segments communicating with each other.

North-South Routing

The routing design considers different levels of routing in the environment, such as number and type of gateways in NSX, dynamic routing protocol, and others.

The following models for north-south traffic exist:

Table 10: Considerations for the Operating Model for North-South Service Routers

North-South Service Router Operating Model	Description	Benefits	Drawbacks
Active-Active	<ul style="list-style-type: none"> Bandwidth independent of the Tier-0 gateway failover model. Configured in active-active equal-cost multi-path (ECMP) mode. Failover takes approximately 2 seconds for virtual edges and is sub-second for bare-metal edges. 	<ul style="list-style-type: none"> Active-active mode can support up to 8 NSX Edge nodes per northbound service router (SR). Availability can be as high as N+7, with up to 8 active-active NSX Edge nodes. Supports ECMP north-south routing on all nodes in the NSX Edge cluster. 	<ul style="list-style-type: none"> Cannot provide some stateful services, such as SNAT or DNAT.
Active-Standby	<ul style="list-style-type: none"> Bandwidth independent of the Tier-0 gateway failover model. Failover takes approximately 2 seconds for virtual edges and is sub-second for bare-metal edges. 	<ul style="list-style-type: none"> Can provide stateful services such as NAT. 	<ul style="list-style-type: none"> Active-standby mode is limited to a single node. Availability is limited to N+1.

BGP North-South Routing for a Single or Multiple Availability Zones

For multiple availability zones, plan for failover of the NSX Edge nodes by configuring BGP so that traffic from the top of rack switches is directed to the first availability zone unless a failure in this zone occurs.

Figure 10: BGP North-South Routing for VMware Cloud Foundation Instances with a Single Availability Zone

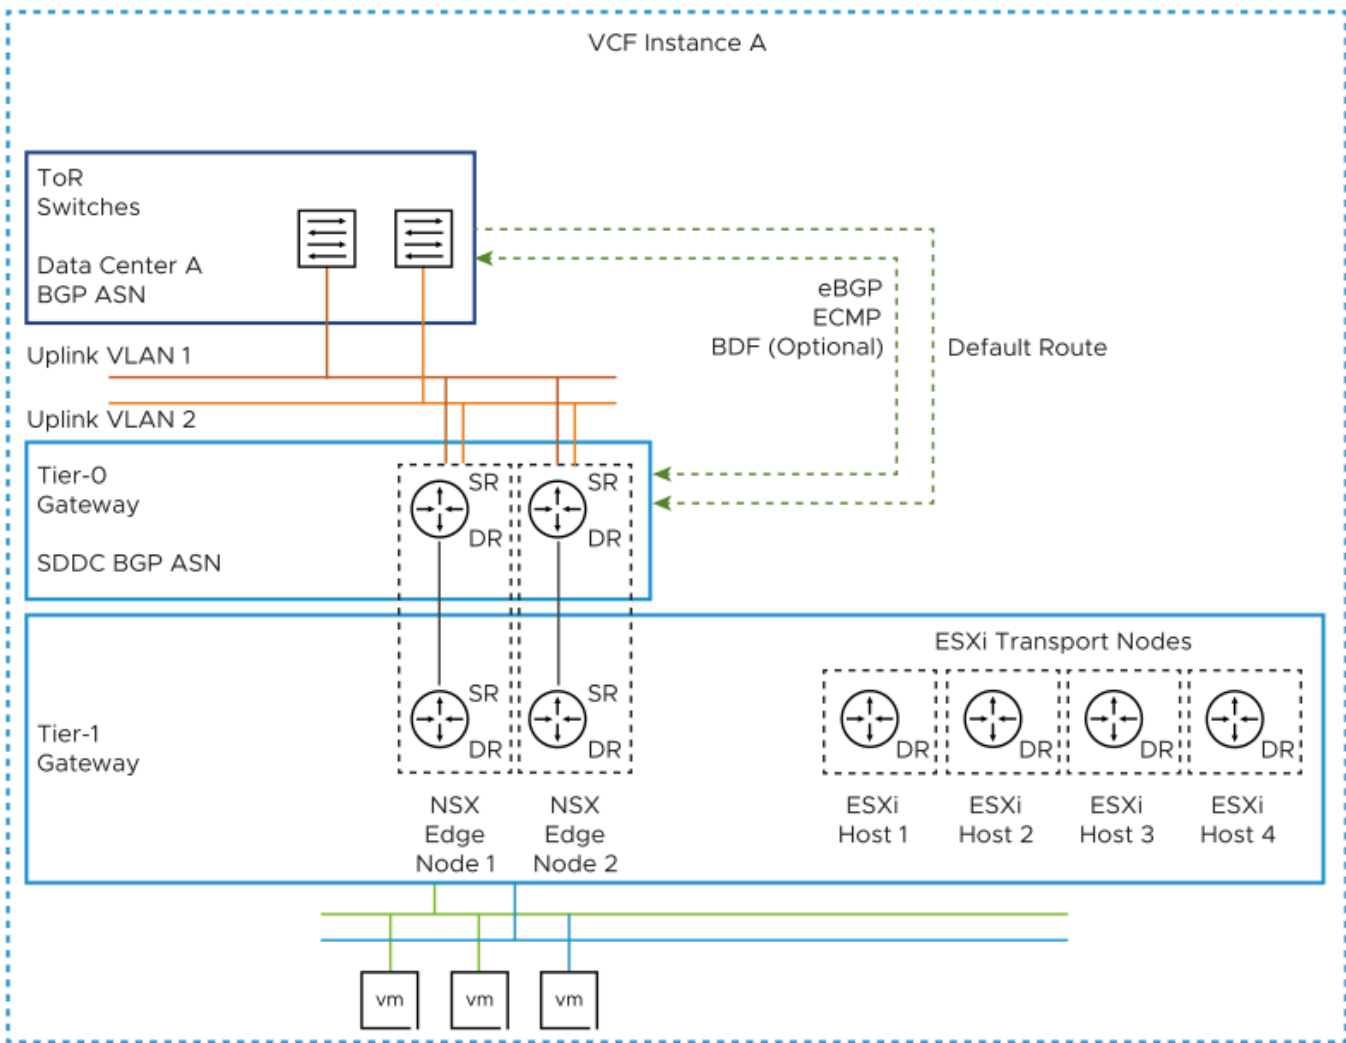
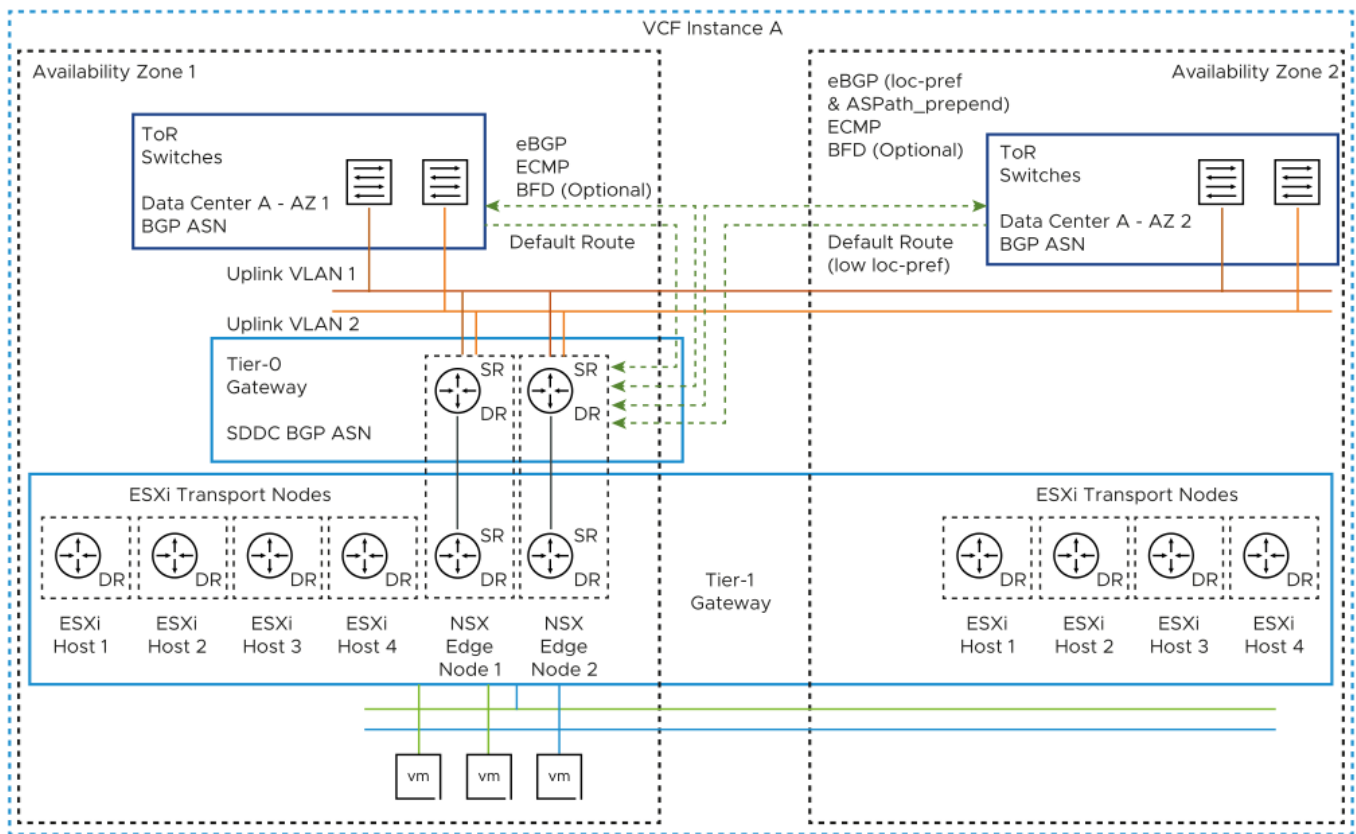


Figure 11: BGP North-South Routing for VMware Cloud Foundation Instances with Multiple Availability Zones

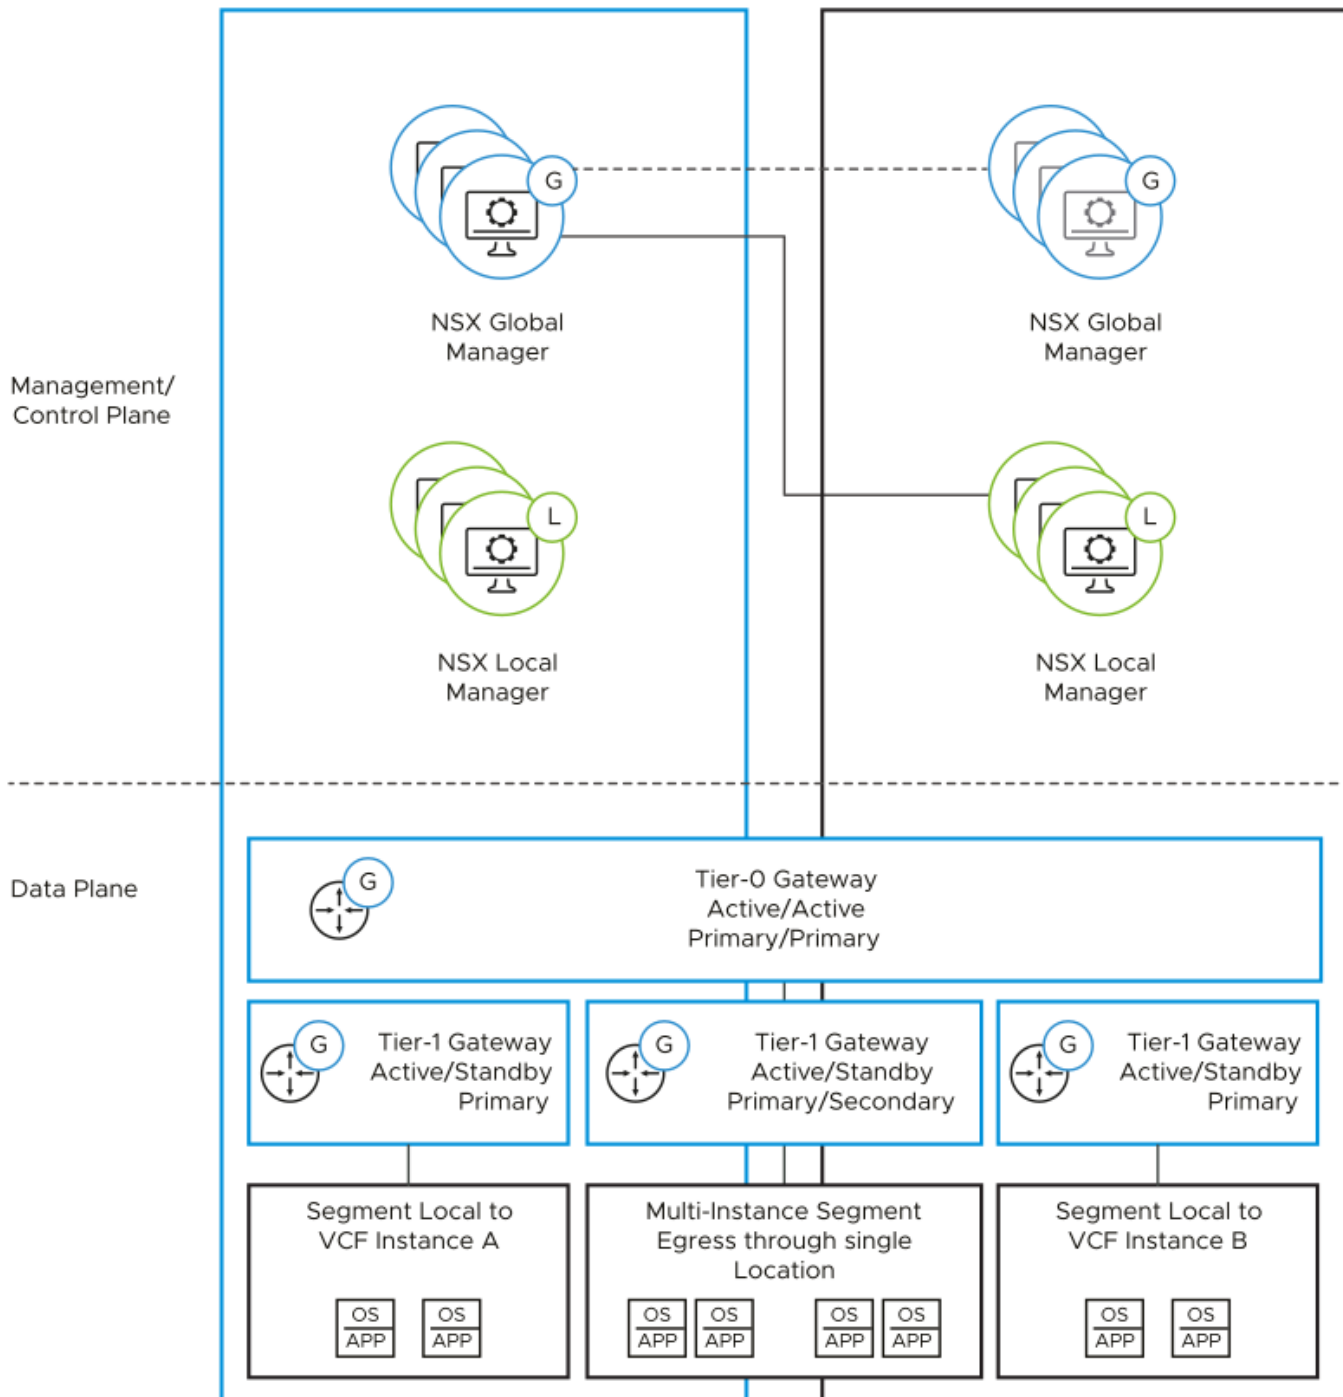


BGP North-South Routing Design for NSX Federation

In a routing design for an environment with VMware Cloud Foundation instances that use NSX Federation, you identify the instances that an SDN network must span and at which physical location ingress and egress traffic should occur.

Local egress allows traffic to leave any location which the network spans. The use of local-egress would require controlling local-ingress to prevent asymmetrical routing. This design does not use local-egress. Instead, this design uses a preferred and failover VMware Cloud Foundation instances for all networks.

Figure 12: BGP North-South Routing for VMware Cloud Foundation Instances with NSX Federation



Tier-0 Gateways with NSX Federation

In NSX Federation, a Tier-0 gateway can span multiple VMware Cloud Foundation instances.

Each VMware Cloud Foundation instance that is in the scope of a Tier-0 gateway can be configured as primary or secondary. A primary instance passes traffic for any other SDN service such as Tier-0 logical segments or Tier-1

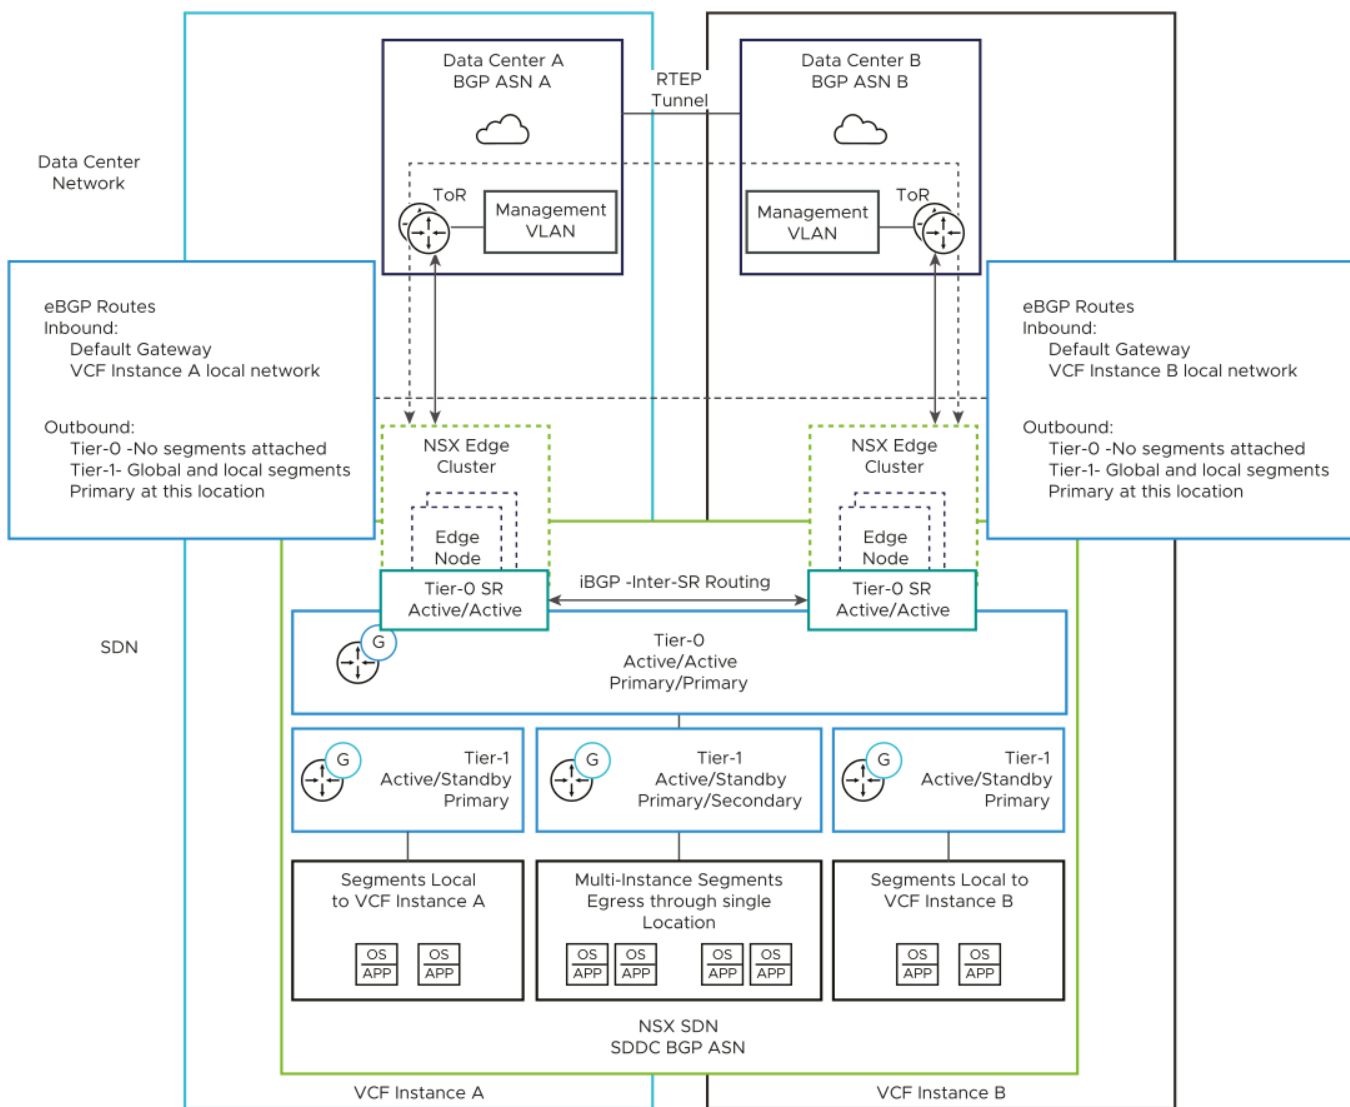
gateways. A secondary instance routes traffic locally but does not egress traffic outside the SDN or advertise networks in the data center.

When deploying an additional VMware Cloud Foundation instance, the Tier-0 gateway in the first instance is extended to the new instance.

In this design, the Tier-0 gateway in each VMware Cloud Foundation instance is configured as primary. Although the Tier-0 gateway technically supports local-egress, the design does not recommend the use of local-egress. Ingress and egress traffic is controlled at the Tier-1 gateway level.

Each VMware Cloud Foundation instance has its own NSX Edge cluster with associated uplink VLANs for north-south traffic flow for that instance. The Tier-0 gateway in each instance peers with the top of rack switches over eBGP.

Figure 13: BGP Peering to Top of Rack Switches for VMware Cloud Foundation Instances with NSX Federation



Tier-1 Gateways with NSX Federation

A Tier-1 gateway can span several VMware Cloud Foundation instances. As with a Tier-0 gateway, you can configure an instance's location as primary or secondary for the Tier-1 gateway. The gateway then passes ingress and egress traffic for the logical segments connected to it.

Any logical segments connected to the Tier-1 gateway follow the span of the Tier-1 gateway. If the Tier-1 gateway spans several VMware Cloud Foundation instances, any segments connected to that gateway become available in both instances.

Using a Tier-1 gateway enables more granular control on logical segments in the first and second V VMware Cloud Foundation instances. You use three Tier-1 gateways - one in each VMware Cloud Foundation instance for segments that are local to the instance, and one for segments which span the two instances.

Table 11: Location Configuration of the Tier-1 Gateways for Multiple VMware Cloud Foundation Instances

Tier-1 Gateway	First VMware Cloud Foundation Instance	Second VMware Cloud Foundation Instance	Ingress and Egress Traffic
Connected to both VMware Cloud Foundation instances	Primary	Secondary	First VMware Cloud Foundation instance Second VMware Cloud Foundation instance
Local to the first VMware Cloud Foundation instance	Primary	-	First VMware Cloud Foundation instance only
Local to the second VMware Cloud Foundation instance	-	Primary	Second VMware Cloud Foundation instance only

The Tier-1 gateway advertises its networks to the connected local-instance unit of the Tier-0 gateway. In the case of primary-secondary location configuration, the Tier-1 gateway advertises its networks only to the Tier-0 gateway unit in the location where the Tier-1 gateway is primary. The Tier-0 gateway unit then re-advertises those networks to the data center in the sites where that Tier-1 gateway is primary. During failover of the components in the first VMware Cloud Foundation instance, an administrator must manually set the Tier-1 gateway in the second VMware Cloud Foundation instance as primary. Then, networks become advertised through the Tier-1 gateway unit in the second instance.

In a Multiple Instance-Multiple Availability Zone topology, the same Tier-0 and Tier-1 gateway architecture applies. The ESXi transport nodes from the second availability zone are also attached to the Tier-1 gateway as per the [nsx-t-routing-for-the-management-domain-in-a-single-region-sddc.dita#FIG_538E0CB8-B373-48BF-90D1-37291E173C77-en](#) design.

BGP Routing Design Requirements and Recommendations for VMware Cloud Foundation

Consider the requirements for the configuration of Tier-0 and Tier-1 gateways for implementing BGP routing in VMware Cloud Foundation, and the best practices for having optimal traffic routing on a standard or stretched cluster in a environment with a single or multiple VMware Cloud Foundation instances.

BGP Routing

The BGP routing design has the following characteristics:

- Enables dynamic routing by using NSX.
- Offers increased scale and flexibility.
- Is a proven protocol that is designed for peering between networks under independent administrative control - data center networks and the NSX SDN.

NOTE

These design recommendations do not include BFD. However, if faster convergence than BGP timers is required, you must enable BFD on the physical network and also on the NSX Tier-0 gateway.

BGP Routing Design Requirements

You must meet the following design requirements for standard and stretched clusters in your routing design for a single VMware Cloud Foundation instance. For NSX Federation, additional requirements exist.

Table 12: BGP Routing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-001	To enable ECMP between the Tier-0 gateway and the Layer 3 devices (ToR switches or upstream devices), create two VLANs. The ToR switches or upstream Layer 3 devices have an SVI on one of the two VLANS, and each Edge node in the cluster has an interface on each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway and provides more resiliency and better bandwidth use in the network.	Additional VLANs are required.
VCF-NSX-BGP-REQD-CFG-002	Assign a named teaming policy to the VLAN segments to the Layer 3 device pair.	Pins the VLAN traffic on each segment to its target edge node interface. From there, the traffic is directed to the host physical NIC that is connected to the target top of rack switch.	None.
VCF-NSX-BGP-REQD-CFG-003	Create a VLAN transport zone for edge uplink traffic.	Enables the configuration of VLAN segments on the N-VDS in the edge nodes.	Additional VLAN transport zones might be required if the edge nodes are not connected to the same top of rack switch pair.
VCF-NSX-BGP-REQD-CFG-004	Deploy a Tier-1 gateway and connect it to the Tier-0 gateway.	Creates a two-tier routing architecture. Abstracts the NSX logical components which interact with the physical data center from the logical components which provide SDN services.	A Tier-1 gateway can only be connected to a single Tier-0 gateway. In cases where multiple Tier-0 gateways are required, you must create multiple Tier-1 gateways.
VCF-NSX-BGP-REQD-CFG-005	Deploy a Tier-1 gateway to the NSX Edge cluster.	Enables stateful services, such as load balancers and NAT, for SDDC management components. Because a Tier-1 gateway always works in active-	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		standby mode, the gateway supports stateful services.	

Table 13: BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-006	Extend the uplink VLANs to the top of rack switches so that the VLANs are stretched between both availability zones.	Because the NSX Edge nodes will fail over between the availability zones, ensures uplink connectivity to the top of rack switches in both availability zones regardless of the zone the NSX Edge nodes are presently in.	You must configure a stretched Layer 2 network between the availability zones by using physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-007	Provide this SVI configuration on the top of the rack switches. <ul style="list-style-type: none"> In the second availability zone, configure the top of rack switches or upstream Layer 3 devices with an SVI on each of the two uplink VLANs. Make the top of rack switch SVI in both availability zones part of a common stretched Layer 2 network between the availability zones. 	Enables the communication of the NSX Edge nodes to the top of rack switches in both availability zones over the same uplink VLANs.	You must configure a stretched Layer 2 network between the availability zones by using the physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-008	Provide this VLAN configuration: <ul style="list-style-type: none"> Use two VLANs to enable ECMP between the Tier-0 gateway and the Layer 3 devices (top of rack switches or Leaf switches). The ToR switches or upstream Layer 3 devices have an SVI to one of the two VLANs and each NSX Edge node has an interface to each VLAN. 	Supports multiple equal-cost routes on the Tier-0 gateway, and provides more resiliency and better bandwidth use in the network.	<ul style="list-style-type: none"> Extra VLANs are required. Requires stretching uplink VLANs between availability zones
VCF-NSX-BGP-REQD-CFG-009	Create an IP prefix list that permits access to route advertisement by any netwo	Used in a route map to prepend a path to one or more autonomous system	You must manually create an IP prefix list that is identical to the default one.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	rk instead of using the default IP prefix list.	(AS-path prepend) for BGP neighbors in the second availability zone.	
VCF-NSX-BGP-REQD-CFG-010	Create a route map-out that contains the custom IP prefix list and an AS-path prepend value set to the Tier-0 local AS added twice.	<ul style="list-style-type: none"> Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. Ensures that all ingress traffic passes through the first availability zone. 	<p>You must manually create the route map.</p> <p>The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.</p>
VCF-NSX-BGP-REQD-CFG-011	Create an IP prefix list that permits access to route advertisement by network 0.0.0.0/0 instead of using the default IP prefix list.	Used in a route map to configure local-reference on learned default-route for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.
VCF-NSX-BGP-REQD-CFG-012	Apply a route map-in that contains the IP prefix list for the default route 0.0.0.0/0 and assign a lower local-preference , for example, 80, to the learned default route and a lower local-preference, for example, 90 any routes learned.	<ul style="list-style-type: none"> Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. Ensures that all egress traffic passes through the first availability zone. 	<p>You must manually create the route map.</p> <p>The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.</p>
VCF-NSX-BGP-REQD-CFG-013	Configure the neighbors of the second availability zone to use the route maps as In and Out filters respectively.	Makes the path in and out of the second availability zone less preferred because the AS path is longer and the local preference is lower. As a result, all traffic passes through the first zone.	The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.

Table 14: BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-014	Extend the Tier-0 gateway to the second VMware Cloud Foundation instance.	<ul style="list-style-type: none"> • Supports ECMP north-south routing on all nodes in the NSX Edge cluster. • Enables support for cross-instance Tier-1 gateways and cross-instance network segments. 	The Tier-0 gateway deployed in the second instance is removed.
VCF-NSX-BGP-REQD-CFG-015	Set the Tier-0 gateway as primary for all VMware Cloud Foundation instances.	<ul style="list-style-type: none"> • In NSX Federation, a Tier-0 gateway lets egress traffic from connected Tier-1 gateways only in its primary locations. • Local ingress and egress traffic is controlled independently at the Tier-1 level. No segments are provisioned directly to the Tier-0 gateway. • A mixture of network spans (local to a VMware Cloud Foundation instance or spanning multiple instances) is enabled without requiring additional Tier-0 gateways and hence edge nodes. • If a failure in a VMware Cloud Foundation instance occurs, the local-instance networking in the other instance remains available without manual intervention. 	None.
VCF-NSX-BGP-REQD-CFG-016	From the global Tier-0 gateway, establish BGP neighbor peering to the ToR switches connected to the second VMware Cloud Foundation instance.	<ul style="list-style-type: none"> • Enables the learning and advertising of routes in the second VMware Cloud Foundation instance. • Facilitates a potential automated failover of networks from the first to the second VMware Cloud Foundation instance. 	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-017	Use a stretched Tier-1 gateway and connect it to the Tier-0 gateway for cross-instance networking.	<ul style="list-style-type: none"> Enables network span between the VMware Cloud Foundation instances because NSX network segments follow the span of the gateway they are attached to. Creates a two-tier routing architecture. 	None.
VCF-NSX-BGP-REQD-CFG-018	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the stretched Tier-1 gateway. Set the first VMware Cloud Foundation instance as primary and the second instance as secondary.	<ul style="list-style-type: none"> Enables cross-instance network span between the first and second VMware Cloud Foundation instances. Enables deterministic ingress and egress traffic for the cross-instance network. If a VMware Cloud Foundation instance failure occurs, enables deterministic failover of the Tier-1 traffic flow. During the recovery of the inaccessible VMware Cloud Foundation instance, enables deterministic failback of the Tier-1 traffic flow, preventing unintended asymmetrical routing. Eliminates the need to use BGP attributes in the first and second VMware Cloud Foundation instances to influence location preference and failover. 	You must manually fail over and fail back the cross-instance network from the standby NSX Global Manager.
VCF-NSX-BGP-REQD-CFG-019	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the local Tier-1 gateway for that VMware Cloud Foundation instance.	<ul style="list-style-type: none"> Enables instance-specific networks to be isolated to their specific instances. Enables deterministic flow of ingress and egress traffic for the instance-specific networks. 	You can use the service router that is created for the Tier-1 gateway for networking services. However, such configuration is not required for network connectivity.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-020	Set each local Tier-1 gateway only as primary in that instance. Avoid setting the gateway as secondary in the other instances.	Prevents the need to use BGP attributes in primary and secondary instances to influence the instance ingress-egress preference.	None.

Table 15: BGP Routing Design Requirements for NSX Multi-Rack Edge Availability for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-MRE-BGP-REQD-CFG-001	To enable ECMP between the Tier-0 gateway and the Layer 3 devices, such as leaf switches or upstream devices, create two separate uplink VLANs for the edge nodes in each rack. The leaf switches or Layer 3 upstream devices have an SVI on one of the two VLANs for each rack, and each edge node in the rack has an interface on each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway and improves resiliency and bandwidth use in the network across multiple racks with a pair of leaf switches in each rack.	Additional VLANs are required.
VCF-NSX-MRE-BGP-REQD-CFG-002	Assign a named teaming policy to the VLAN segments to the Layer 3 device pair for each rack.	Pins the VLAN traffic on each segment to its target edge node interface. From there, the traffic is directed to the host physical NIC that is connected to the target leaf switch in the rack.	None.

BGP Routing Design Recommendations

In your routing design for a single VMware Cloud Foundation instance, you can apply certain best practices for standard and stretched clusters. For NSX Federation, additional recommendations are available.

Table 16: BGP Routing Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
VCF-NSX-BGP-RCMD-CFG-001	Deploy an active-active Tier-0 gateway.	Supports ECMP north-south routing on all Edge nodes in the NSX Edge cluster.	Active-active Tier-0 gateways cannot provide stateful services such as NAT.
VCF-NSX-BGP-RCMD-CFG-002	Configure the BGP Keep Alive Timer to 4 and Hold Down Timer to 12 or lower between the top of tack	Provides a balance between failure detection between the top of rack switches and the Tier-0 gateway, and	By using longer timers to detect if a router is not responding, the data about such a router remains in the

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
	switches and the Tier-0 gateway.	overburdening the top of rack switches with keep-alive traffic.	routing table longer. As a result, the active router continues to send traffic to a router that is down. These timers must be aligned with the data center fabric design of your organization.
VCF-NSX-BGP-RCMD-CFG-003	Do not enable Graceful Restart between BGP neighbors.	Avoids loss of traffic. On the Tier-0 gateway, BGP peers from all the gateways are always active. On a failover, the Graceful Restart capability increases the time a remote neighbor takes to select an alternate Tier-0 gateway. As a result, BFD-based convergence is delayed.	None.
VCF-NSX-BGP-RCMD-CFG-004	Enable helper mode for Graceful Restart mode between BGP neighbors.	Avoids loss of traffic. During a router restart, helper mode works with the graceful restart capability of upstream routers to maintain the forwarding table which in turn will forward packets to a down neighbor even after the BGP timers have expired causing loss of traffic.	None.
VCF-NSX-BGP-RCMD-CFG-005	Enable Inter-SR iBGP routing.	In the event that an edge node has all of its northbound eBGP sessions down, north-south traffic will continue to flow by routing traffic to a different edge node.	None.
VCF-NSX-BGP-RCMD-CFG-006	Deploy a Tier-1 gateway in non-preemptive failover mode.	Ensures that after a failed NSX Edge transport node is back online, it does not take over the gateway services thus preventing a short service outage.	None.
VCF-NSX-BGP-RCMD-CFG-007	Enable standby relocation of the Tier-1 gateway.	Ensures that if an edge failure occurs, a standby Tier-1 gateway is created on another edge node.	None.

Table 17: BGP Routing Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-BGP-RCMD-CFG-008	Use Tier-1 gateways to control the span of networks and ingress and egress traffic in the VMware Cloud Foundation instances.	Enables a mixture of network spans (isolated to a VMware Cloud Foundation instance or spanning multiple instances) without requiring additional Tier-0 gateways and hence edge nodes.	To control location span, a Tier-1 gateway must be assigned to an edge cluster and hence has the Tier-1 SR component. East-west traffic between Tier-1 gateways with SRs need to physically traverse an edge node.
VCF-NSX-BGP-RCMD-CFG-009	Allocate a Tier-1 gateway in each instance for instance-specific networks and connect it to the stretched Tier-0 gateway.	<ul style="list-style-type: none"> Creates a two-tier routing architecture. Enables local-instance networks that are not to span between the VMware Cloud Foundation instances. Guarantees that local-instance networks remain available if a failure occurs in another VMware Cloud Foundation instance. 	None.

Lifecycle Management Design for VMware Cloud Foundation

In a VMware Cloud Foundation instance, you use SDDC Manager for lifecycle management of the management components in the entire instance except for NSX Global Manager and VMware Aria Suite Lifecycle. VMware Aria Suite Lifecycle manages the lifecycle of the components that it deploys.

Lifecycle management of a VMware Cloud Foundation instance is the process of performing patch updates or upgrades to the underlying management components.

Table 18: Lifecycle Management for VMware Cloud Foundation

Component	Management Domain	VI Workload Domain
SDDC Manager	SDDC Manager performs its own life cycle management.	Not applicable
NSX Local Manager	SDDC Manager uses the NSX upgrade coordinator service in the NSX Local Manager.	
NSX Edges	SDDC Manager uses the NSX upgrade coordinator service in NSX Manager.	
NSX Global Manager	You manually use the NSX upgrade coordinator service in the NSX Global Manager.	
vCenter Server	You use SDDC Manager for life cycle management of all vCenter Server instances.	
ESXi	<ul style="list-style-type: none"> SDDC Manager uses vSphere Lifecycle Manager baselines or vSphere Lifecycle Manager images to update and upgrade the ESXi hosts. Custom vendor ISOs are supported and might be required according to the ESXi hardware in use. 	

Table continued on next page

Continued from previous page

Component	Management Domain	VI Workload Domain
VMware Aria Suite Lifecycle	VMware Aria Suite Lifecycle performs its own life cycle management.	Not applicable

VMware Cloud Foundation Lifecycle Management Requirements

Consider the design requirements for automated and centralized lifecycle management in the context of the entire VMware Cloud Foundation environment.

Table 19: Lifecycle Management Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-001	Use SDDC Manager to perform the life cycle management of the following components: <ul style="list-style-type: none"> • SDDC Manager • NSX Manager • NSX Edges • vCenter Server • ESXi 	Because the deployment scope of SDDC Manager covers the full VMware Cloud Foundation stack, SDDC Manager performs patching, update, or upgrade of these components across all workload domains.	The operations team must understand and be aware of the impact of a patch, update, or upgrade operation by using SDDC Manager.
VCF-LCM-REQD-002	Use VMware Aria Suite Lifecycle to manage the life cycle of the following components: <ul style="list-style-type: none"> • VMware Aria Suite Lifecycle • Workspace ONE Access 	VMware Aria Suite Lifecycle automates the life cycle of VMware Aria Suite Lifecycle and Workspace ONE Access.	<ul style="list-style-type: none"> • You must deploy VMware Aria Suite Lifecycle by using SDDC Manager. • You must manually apply Workspace ONE Access patches, updates, and hotfixes. Patches, updates, and hotfixes for Workspace ONE Access are not generally managed by VMware Aria Suite Lifecycle.

Table 20: Lifecycle Management Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-003	Use the upgrade coordinator in NSX to perform life cycle management on the NSX Global Manager appliances.	The version of SDDC Manager in this design is not currently capable of life cycle operations (patching, update, or upgrade) for NSX Global Manager.	<ul style="list-style-type: none"> • You must explicitly plan upgrades of the NSX Global Manager nodes. An upgrade of the NSX Global Manager nodes might require a cascading upgrade of the NSX Local Manager nodes and underlying SDDC Manager infrastructure

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
			before upgrading the NSX Global Manager nodes. <ul style="list-style-type: none"> You must always align the version of the NSX Global Manager nodes with the rest of the SDDC stack in VMware Cloud Foundation.
VCF-LCM-REQD-004	Establish an operations practice to ensure that prior to the upgrade of any workload domain, the impact of any version upgrades is evaluated in relation to the need to upgrade NSX Global Manager.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.
VCF-LCM-REQD-005	Establish an operations practice to ensure that prior to the upgrade of the NSX Global Manager, the impact of any version change is evaluated against the existing NSX Local Manager nodes and workload domains.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.

Logging and Monitoring Design for VMware Cloud Foundation

By using VMware or third-party components, collect log data from all SDDC management components in your VMware Cloud Foundation environment in a central place. You can use VMware Aria Operations for Logs as the central platform because of its native integration with VMware Aria Suite Lifecycle.

After you deploy VMware Aria Operations for Logs by using VMware Aria Suite Lifecycle in VMware Cloud Foundation mode, SDDC Manager configures VMware Aria Suite Lifecycle logging to VMware Aria Operations for Logs over the log ingestion API. For information about on-premises VMware Aria Operations for Logs in VMware Cloud Foundation, see [Intelligent Logging and Analytics for VMware Cloud Foundation](#).

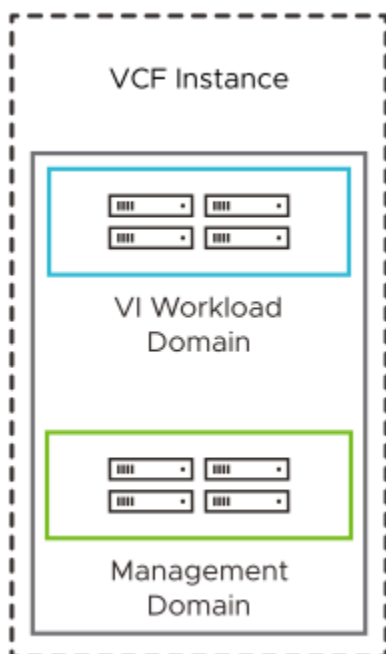
VMware Cloud Foundation Topology Design Blueprints

A *VMware Cloud Foundation* topology design blueprint is a collection of design requirements and recommendations based on a chosen architecture model, workload domain type, and topology. It can be used as a full end-to-end design for a *VMware Cloud Foundation* deployment.

Topology Design Blueprint One: Single Instance - Single Availability Zone

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology which includes one *VMware Cloud Foundation* instance, with a management domain and one or more VI workload domains in a single availability zone for an organization called Rainpole.

Figure 14: Single VMware Cloud Foundation Instance - Single Availability Zone



Design Choices for Design Blueprint One

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 21: Design Choices for Design Blueprint One

Design Aspect	Choice Made
Architecture model	Standard
Workload domain type	Management domain and VI workload domains
Topology	Single Instance - Single Availability Zone
Physical network configuration	Leaf-Spine
Routing configuration	BGP
Workload domain principal storage	vSAN
VMware Aria Suite Lifecycle	Included
Workspace ONE Access	Standard Workspace ONE Access

Design Elements for Design Blueprint One**Table 22: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 23: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations

Table 24: Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Recommendations
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Recommendations
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendation
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Recommendations
Routing	BGP Routing Design Requirements
	BGP Routing Design Recommendations
Overlay	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network	Application Virtual Network Design Requirements
Load balancing	Load Balancing Design Requirements
SDDC Manager	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 25: VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle	VMware Aria Suite Lifecycle Design Requirements

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Recommendations

Table 26: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 27: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

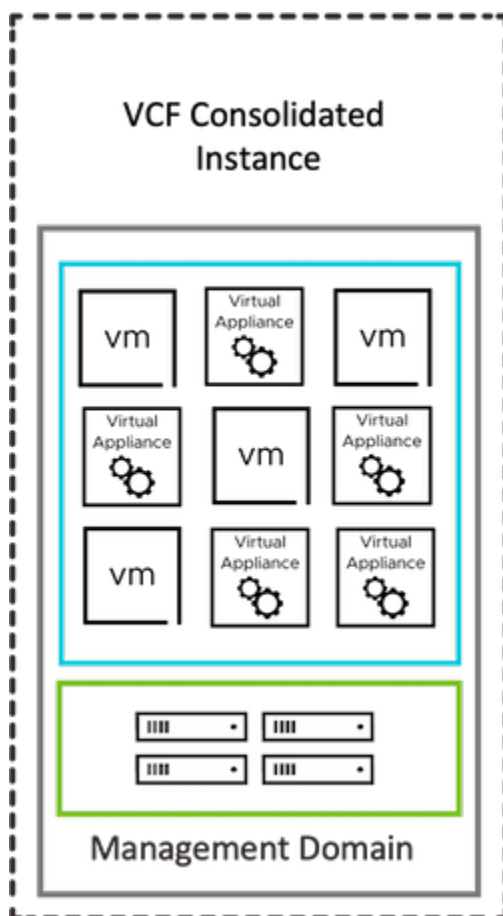
Table 28: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

Topology Design Blueprint Two: Consolidated Single Instance - Single Availability Zone

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology which includes one VMware Cloud Foundation instance where the management domain runs both management and customer workloads in a single availability zone for an organization called Rainpole.

Figure 15: Consolidated VMware Cloud Foundation Instance - Single Availability Zone

**Design Choices for Design Blueprint Two**

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 29: Design Choices for Design Blueprint Two

Design Aspect	Choice Made
Architecture model	Consolidated
Workload domain type	Consolidated
Topology	Single Instance - Single Availability Zone
Physical network configuration	Leaf-Spine
Routing configuration	BGP
Workload domain principal storage	vSAN
VMware Aria Suite Lifecycle	Included
Workspace ONE Access	Standard Workspace ONE Access

Design Elements for Design Blueprint Two**Table 30: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 31: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations

Table 32: Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Recommendations
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Recommendations
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendation
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Recommendations
Routing	BGP Routing Design Requirements
	BGP Routing Design Recommendations
Overlay	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network	Application Virtual Network Design Requirements
Load balancing	Load Balancing Design Requirements
SDDC Manager	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 33: VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle	VMware Aria Suite Lifecycle Design Requirements

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Recommendations

Table 34: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 35: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

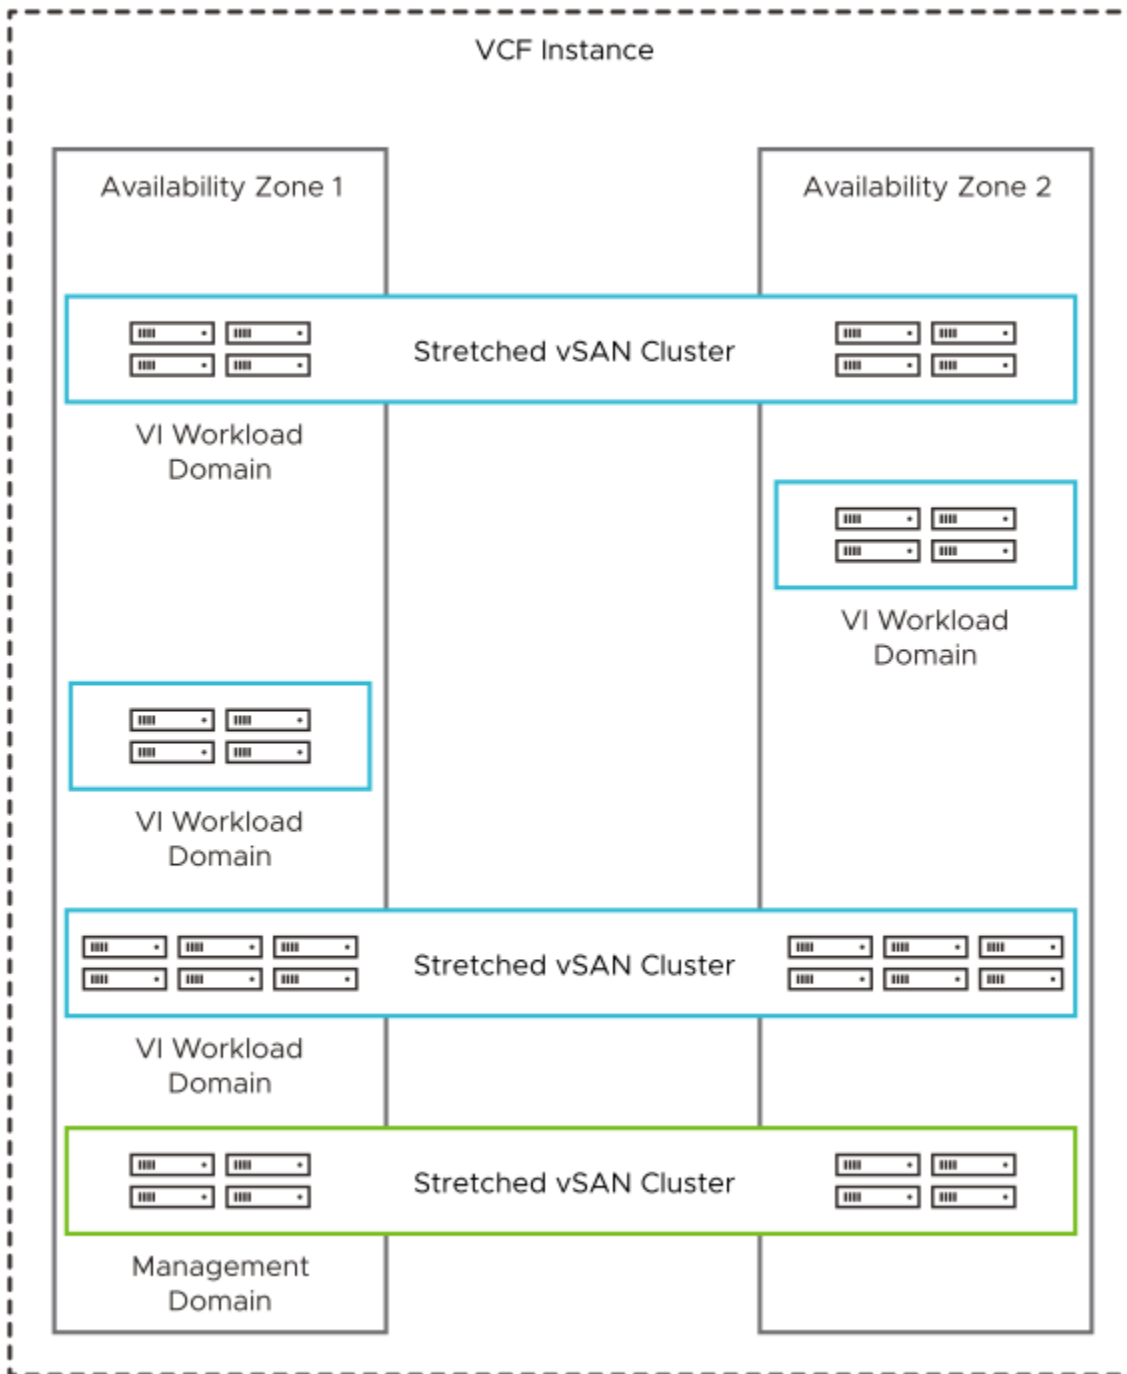
Table 36: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

Topology Design Blueprint Three: Single Instance - Multiple Availability Zones

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology that includes one VMware Cloud Foundation instance with multiple availability zones for an organization called Rainpole.

Figure 16: Single VMware Cloud Foundation Instance - Multiple Availability Zones



Design Choices for Design Blueprint Three

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 37: Design Choices for Design Blueprint Three

Design Aspect	Choice Made
Architecture model	Standard
Workload domain type	Management domain and VI workload domains
Topology	Single Instance - Multiple Availability Zones per instance
Physical network configuration	Leaf-Spine
Routing configuration	BGP
Workload domain principal storage	vSAN
VMware Aria Suite Lifecycle	Included
Workspace ONE Access	Standard Workspace ONE Access

Design Elements for Design Blueprint Three**Table 38: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 39: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations
	Leaf-Spine Physical Network Design Recommendations for Stretched Clusters

Table 40: Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere cluster	vSphere Cluster Design Requirements

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendation
	NSX Manager Design Recommendations for Stretched Clusters
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Recommendations
Overlay	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network	Application Virtual Network Design Requirements
Load balancing	Load Balancing Design Requirements
SDDC Manager	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 41: VI Workload Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single SSO Domain Topology
vSphere cluster	vSphere Cluster Design Requirements VMware Cloud Foundation

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendations
	NSX Manager Design Recommendations for Stretched Clusters
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Recommendations
Overlay	Overlay Design Requirements
	Overlay Design Recommendations

Table 42: VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle	VMware Aria Suite Lifecycle Design Requirements
	VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Requirements for Stretched Clusters
	Workspace ONE Access Design Recommendations

Table 43: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 44: Account and Password Management Design Elements

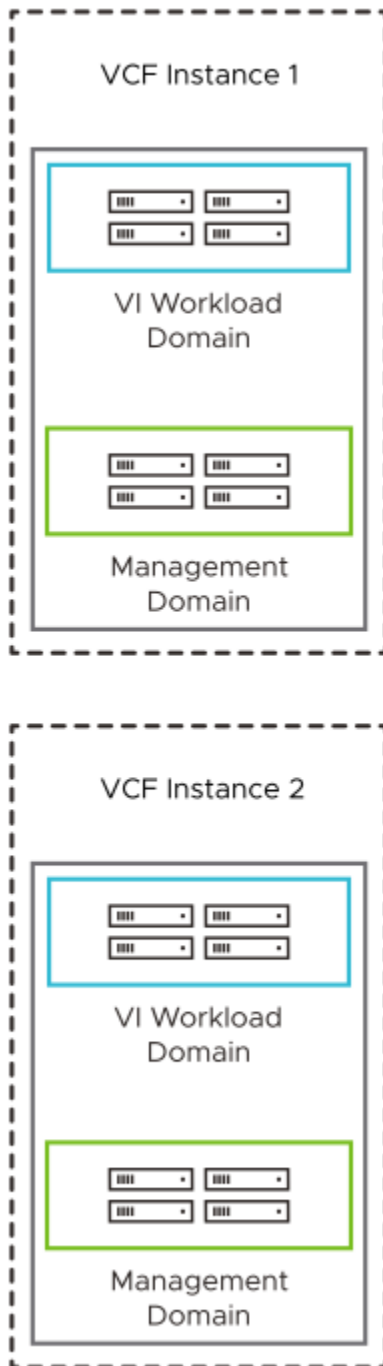
Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

Table 45: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

Topology Design Blueprint Four: Multiple Instance - Single Availability Zone

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology that includes multiple VMware Cloud Foundation instances, each instance containing a single availability zone, for an organization called Rainpole.

Figure 17: Multiple VMware Cloud Foundation Instances - Single Availability Zone per Instance**Design Choices for Design Blueprint Four**

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 46: Design Choices for Design Blueprint Four

Design Aspect	Choice Made
Architecture model	Standard
Workload domain type	Management domain and VI workload domains
Topology	Multiple Instances - Single Availability Zone per instance
Physical network configuration	Leaf-Spine
NSX Federation configuration	NSX Federation between VCF Instances
Routing configuration	BGP
Workload domain principal storage	vSAN
VMware Aria Suite Lifecycle	Included
Workspace ONE Access	Standard Workspace ONE Access

Design Elements for Design Blueprint Four**Table 47: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 48: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Requirements for NSX Federation
	Leaf-Spine Physical Network Design Recommendations
	Leaf-Spine Physical Network Design Recommendations for Stretched Clusters
	Leaf-Spine Physical Network Design Recommendations for NSX Federation

Table 49: Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendation
	NSX Manager Design Recommendations for Stretched Clusters
NSX Global Manager	NSX Global Manager Design Requirements for NSX Federation
	NSX Global Manager Design Recommendations for NSX Federation
	NSX Global Manager Design Recommendations for Stretched Clusters
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Requirements for NSX Federation
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Requirements for NSX Federation
	BGP Routing Design Recommendations
	BGP Routing Design Recommendations for NSX Federation
Overlay	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network	Application Virtual Network Design Requirements
	Application Virtual Network Design Requirements for NSX Federation
Load balancing	Load Balancing Design Requirements
	Load Balancing Design Requirements for NSX Federation
SDDC Manager	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 50: VI Workload Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single SSO Domain Topology
vSphere cluster	vSphere Cluster Design Requirements VMware Cloud Foundation
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendations
	NSX Manager Design Recommendations for Stretched Clusters
NSX Global Manager	NSX Global Manager Design Requirements for NSX Federation
	NSX Global Manager Design Recommendations for NSX Federation
	NSX Global Manager Design Recommendations for Stretched Clusters
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Requirements for NSX Federation
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Requirements for NSX Federation
	BGP Routing Design Recommendations
	BGP Routing Design Recommendations for NSX Federation

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
Overlay	Overlay Design Requirements
	Overlay Design Recommendations

Table 51: VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle	VMware Aria Suite Lifecycle Design Requirements
	VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters
	VMware Aria Suite Lifecycle Design Requirements for NSX Federation
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Requirements for Stretched Clusters
	Workspace ONE Access Design Requirements for NSX Federation
	Workspace ONE Access Design Recommendations

Table 52: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 53: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

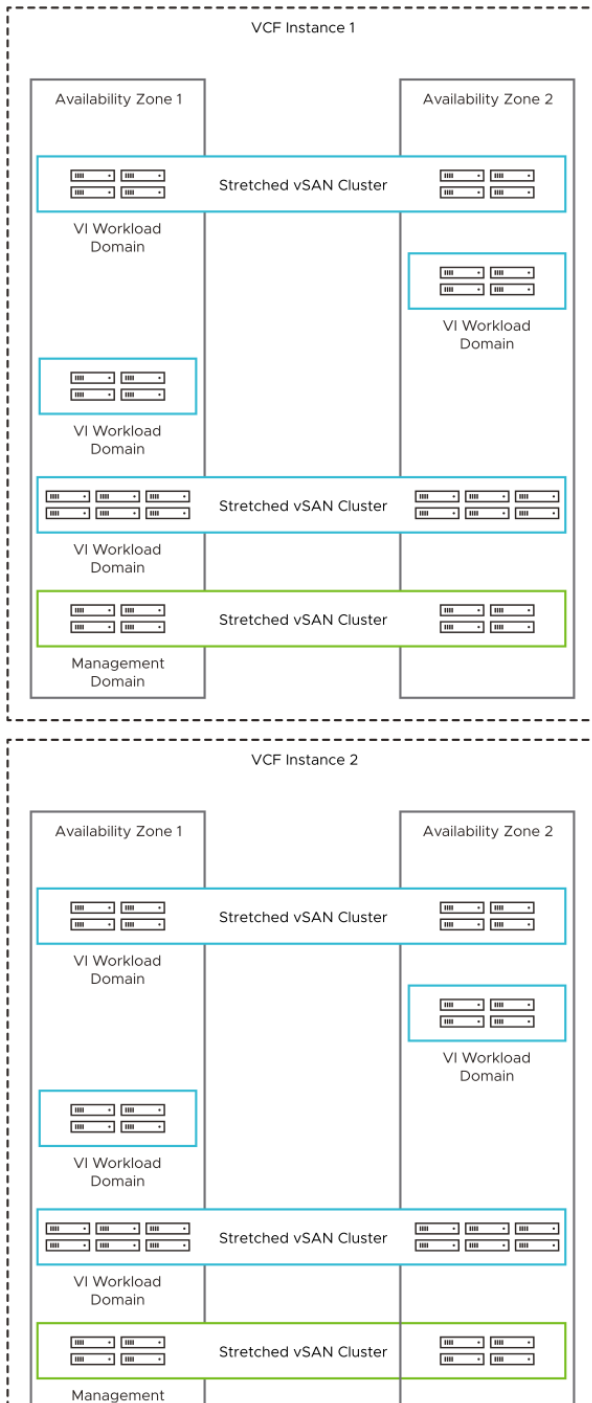
Table 54: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

Topology Design Blueprint Five: Multiple Instance - Multiple Availability Zones

This design blueprint lists the design choices and resulting requirements and recommendations to set up a topology that includes multiple VMware Cloud Foundation instances, each instance containing multiple availability zones, for an organization called Rainpole.

Figure 18: Multiple VMware Cloud Foundation Instances - Multiple Availability Zones per Instance



Design Choices for Design Blueprint Five

Rainpole has made the following choices for its VMware Cloud Foundation deployment:

Table 55: Design Choices for Design Blueprint Five

Design Aspect	Choice Made
Architecture model	Standard
Workload domain type	Management domain and VI workload domains
Topology	Multiple Instances - Multiple Availability Zones per instance
Physical network configuration	Leaf-Spine
Routing configuration	BGP
NSX Federation configuration	NSX Federation between VCF Instances
Workload domain principal storage	vSAN
VMware Aria Suite Lifecycle	Included
Workspace ONE Access	Standard Workspace ONE Access

Design Elements for Design Blueprint Five**Table 56: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 57: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Requirements for NSX Federation
	Leaf-Spine Physical Network Design Recommendations
	Leaf-Spine Physical Network Design Recommendations for Stretched Clusters
	Leaf-Spine Physical Network Design Recommendations for NSX Federation

Table 58: Management Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single vCenter Single Sign-On Domain Topology
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendation
	NSX Manager Design Recommendations for Stretched Clusters
NSX Global Manager	NSX Global Manager Design Requirements for NSX Federation
	NSX Global Manager Design Recommendations for NSX Federation
	NSX Global Manager Design Recommendations for Stretched Clusters
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Requirements for NSX Federation
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Requirements for NSX Federation
	BGP Routing Design Recommendations
	BGP Routing Design Recommendations for NSX Federation
Overlay	Overlay Design Requirements
	Overlay Design Recommendations
Application Virtual Network	Application Virtual Network Design Requirements
	Application Virtual Network Design Requirements for NSX Federation
Load balancing	Load Balancing Design Requirements
	Load Balancing Design Requirements for NSX Federation
SDDC Manager	SDDC Manager Design Requirements
	SDDC Manager Design Recommendations

Table 59: VI Workload Domain Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Stretched Clusters
	vSAN Design Recommendations
	vSAN Design Recommendations for Stretched Clusters
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vCenter Server	vCenter Server Design Requirements
	vCenter Server Design Recommendations
	vCenter Server Design Recommendations for Stretched Clusters
vCenter Single Sign-On	vCenter Single Sign-on Design Requirements for Multiple vCenter - Single SSO Domain Topology
vSphere cluster	vSphere Cluster Design Requirements VMware Cloud Foundation
	vSphere Cluster Design Requirements for Stretched Clusters
	vSphere Cluster Design Recommendations
	vSphere Cluster Design Recommendations for Stretched Clusters
vSphere networking	vSphere Networking Design Recommendations
NSX Manager	NSX Manager Design Requirements
	NSX Manager Design Recommendations
	NSX Manager Design Recommendations for Stretched Clusters
NSX Global Manager	NSX Global Manager Design Requirements for NSX Federation
	NSX Global Manager Design Recommendations for NSX Federation
	NSX Global Manager Design Recommendations for Stretched Clusters
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Requirements for NSX Federation
	NSX Edge Design Recommendations
	NSX Edge Design Recommendations for Stretched Clusters
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Stretched Clusters
	BGP Routing Design Requirements for NSX Federation
	BGP Routing Design Recommendations
	BGP Routing Design Recommendations for NSX Federation

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
Overlay	Overlay Design Requirements
	Overlay Design Recommendations

Table 60: VMware Aria Suite Lifecycle and Workspace ONE Access Design Elements

Design Area	Applicable Design Elements
VMware Aria Suite Lifecycle	VMware Aria Suite Lifecycle Design Requirements
	VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters
	VMware Aria Suite Lifecycle Design Requirements for NSX Federation
	VMware Aria Suite Lifecycle Design Recommendations
Workspace ONE Access	Workspace ONE Access Design Requirements
	Workspace ONE Access Design Requirements for Stretched Clusters
	Workspace ONE Access Design Requirements for NSX Federation
	Workspace ONE Access Design Recommendations

Table 61: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 62: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

Table 63: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

VMware Cloud Foundation vSphere Cluster Design Patterns

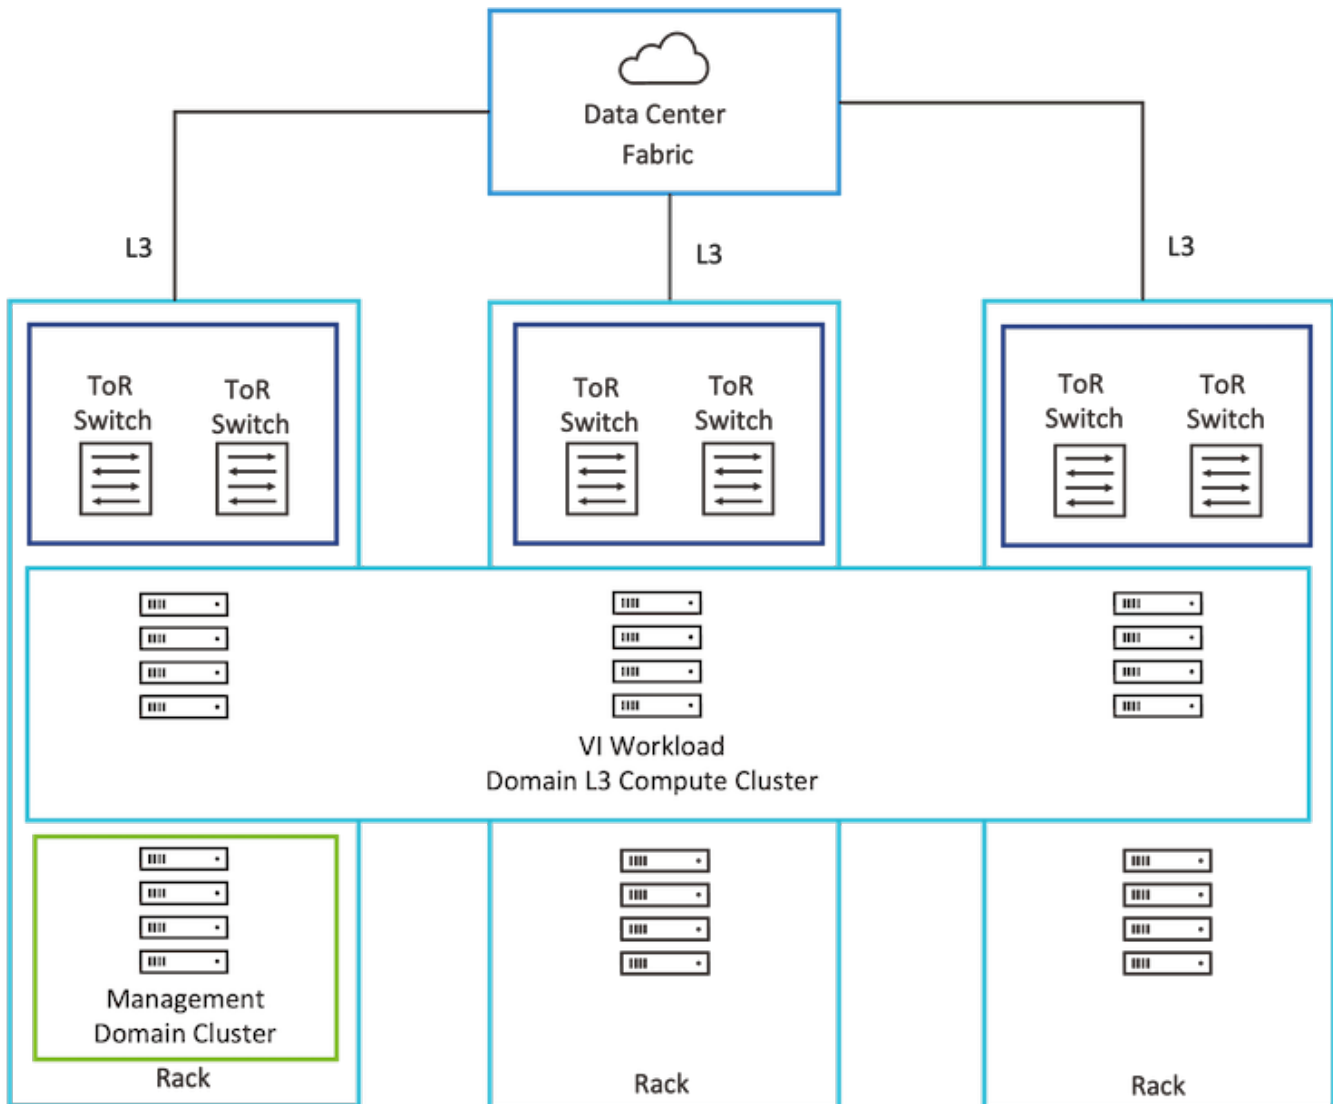
A VMware Cloud Foundation cluster design pattern is a collection of design requirements and recommendations based on a specific cluster design. A cluster design pattern is based on a chosen architecture model, workload domain type, and topology. It covers a particular aspect of a design for a VMware Cloud Foundation deployment rather than all design decisions for a full topology.

vSphere Cluster Design Pattern One: Multi-Rack Compute VI Workload Domain Cluster

This design pattern lists the design choices and resulting requirements and recommendations to deploy a compute cluster that spans multiple racks with vSAN storage and a Layer 3 network fabric in a VI workload domain.

This design provides rack resiliency for the applications running on the cluster that span multiple racks in a VMware Cloud Foundation instance with a single availability zone.

Figure 19: Multi-Rack Compute Cluster with Layer 3 Networks



Design Choices for vSphere Cluster Design Pattern One

This design requires the following VMware Cloud Foundation deployment choices:

Table 64: Design Choices for vSphere Cluster Design Pattern One

Design Aspect	Choice Made
Architecture model	Standard
Workload domain type	VI workload domain
Topology	Single Instance - Single Availability Zone
Workload domain cluster to rack mapping	Workload domain cluster spanning multiple racks
Physical network configuration	Leaf-spine
Workload domain principal storage	vSAN

Design Elements for vSphere Cluster Design Pattern One**Table 65: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 66: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Requirements for L3 Multi-Rack Compute Cluster
	Leaf-Spine Physical Network Design Recommendations

Table 67: Multi-Rack Compute VI Workload Domain Cluster Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Requirements for Multi-Rack Compute VI Workload Domain Cluster
	vSAN ESA Design Requirements for Multi-Rack Compute VI Workload Domain Cluster (applicable for vSAN ESA)
	vSAN OSA Design Requirements for Multi-Rack Compute VI Workload Domain Cluster (applicable for vSAN OSA)
	vSAN Design Recommendations
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Requirements for Multi-Rack Compute VI Workload Domain Cluster
	vSphere Cluster Design Recommendations
vSphere networking	vSphere Networking Design Requirements for Multi-Rack Compute VI Workload Domain Cluster

Table continued on next page

Continued from previous page

Design Area	Applicable Design Elements
Overlay	vSphere Networking Design Recommendations
	Overlay Design Requirements
	Overlay Design Requirements for Multi-Rack Compute V1 Workload Domain Cluster
	Overlay Design Recommendations

Table 68: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 69: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

Table 70: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

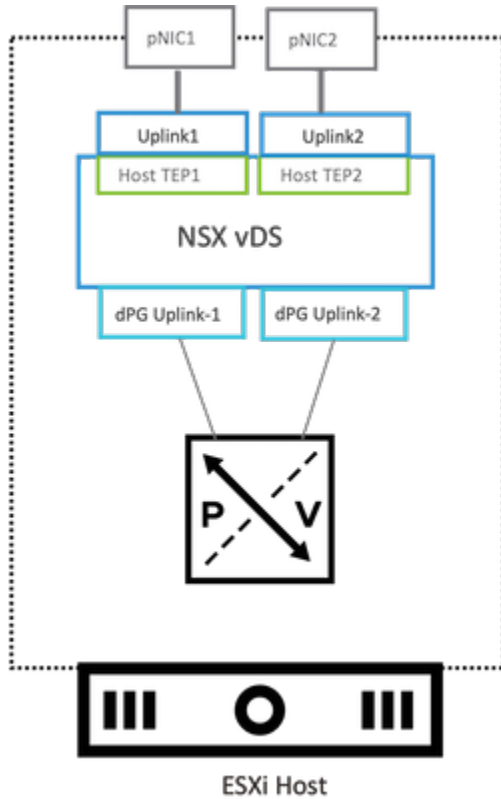
VMware Cloud Foundation NSX Edge Cluster Design Patterns

A VMware Cloud Foundation NSX Edge cluster design pattern is a collection of design requirements and recommendations based on a specific NSX Edge design. Designs are based on availability and performance requirements but also manageability from SDDC Manager. An NSX Edge cluster design pattern covers a particular aspect of a design for a VMware Cloud Foundation required for an NSX Edge cluster deployment rather than all design decisions for a full topology.

NSX Edge Cluster Design Pattern One: Dedicated Edge Scale and Performance

This design pattern lists the design choices and resulting requirements and recommendations to deploy an NSX Edge cluster on a dedicated vSphere cluster with a specific configuration to achieve maximum performance with consistent bandwidth allocated to edge traffic. This design can also be combined with multi-rack edge availability to provide rack resiliency and a high level of availability for the edge cluster with edge VMs in both racks in the VMware Cloud Foundation instance.

Figure 20: Dedicated Edge Scale and Performance



Design Choices for Edge Cluster Design Pattern One

This design requires the following VMware Cloud Foundation deployment choices:

Table 71: Design Choices for NSX Edge Cluster Design Pattern One

Design Aspect	Choice Made
Architecture model	Standard
Workload domain type	VI workload domain
Topology	Multiple Instances - Single Availability Zone per Instance Single Instance - Single Availability Zone (Only supported topology if combining with Multi-Rack Edge Availability)
Workload Domain Cluster to Rack Mapping	Workload domain cluster per Rack
Physical network configuration	Leaf-spine
Workload domain principal storage	vSAN

Design Elements for NSX Edge Cluster Design Pattern One**Table 72: External Services Design Elements**

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 73: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations
	Leaf-Spine Physical Network Design Recommendations for Dedicated Edge Scale and Performance

Table 74: Dedicated Edge Scale and Performance Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Recommendations
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Recommendations
vSphere networking	vSphere Networking Design Requirements for Dedicated Edge Scale and Performance
	vSphere Networking Design Recommendations
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Recommendations
Routing	BGP Routing Design Requirements
	BGP Routing Design Recommendations
Overlay	Overlay Design Requirements
	Overlay Design Recommendations

Table 75: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 76: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

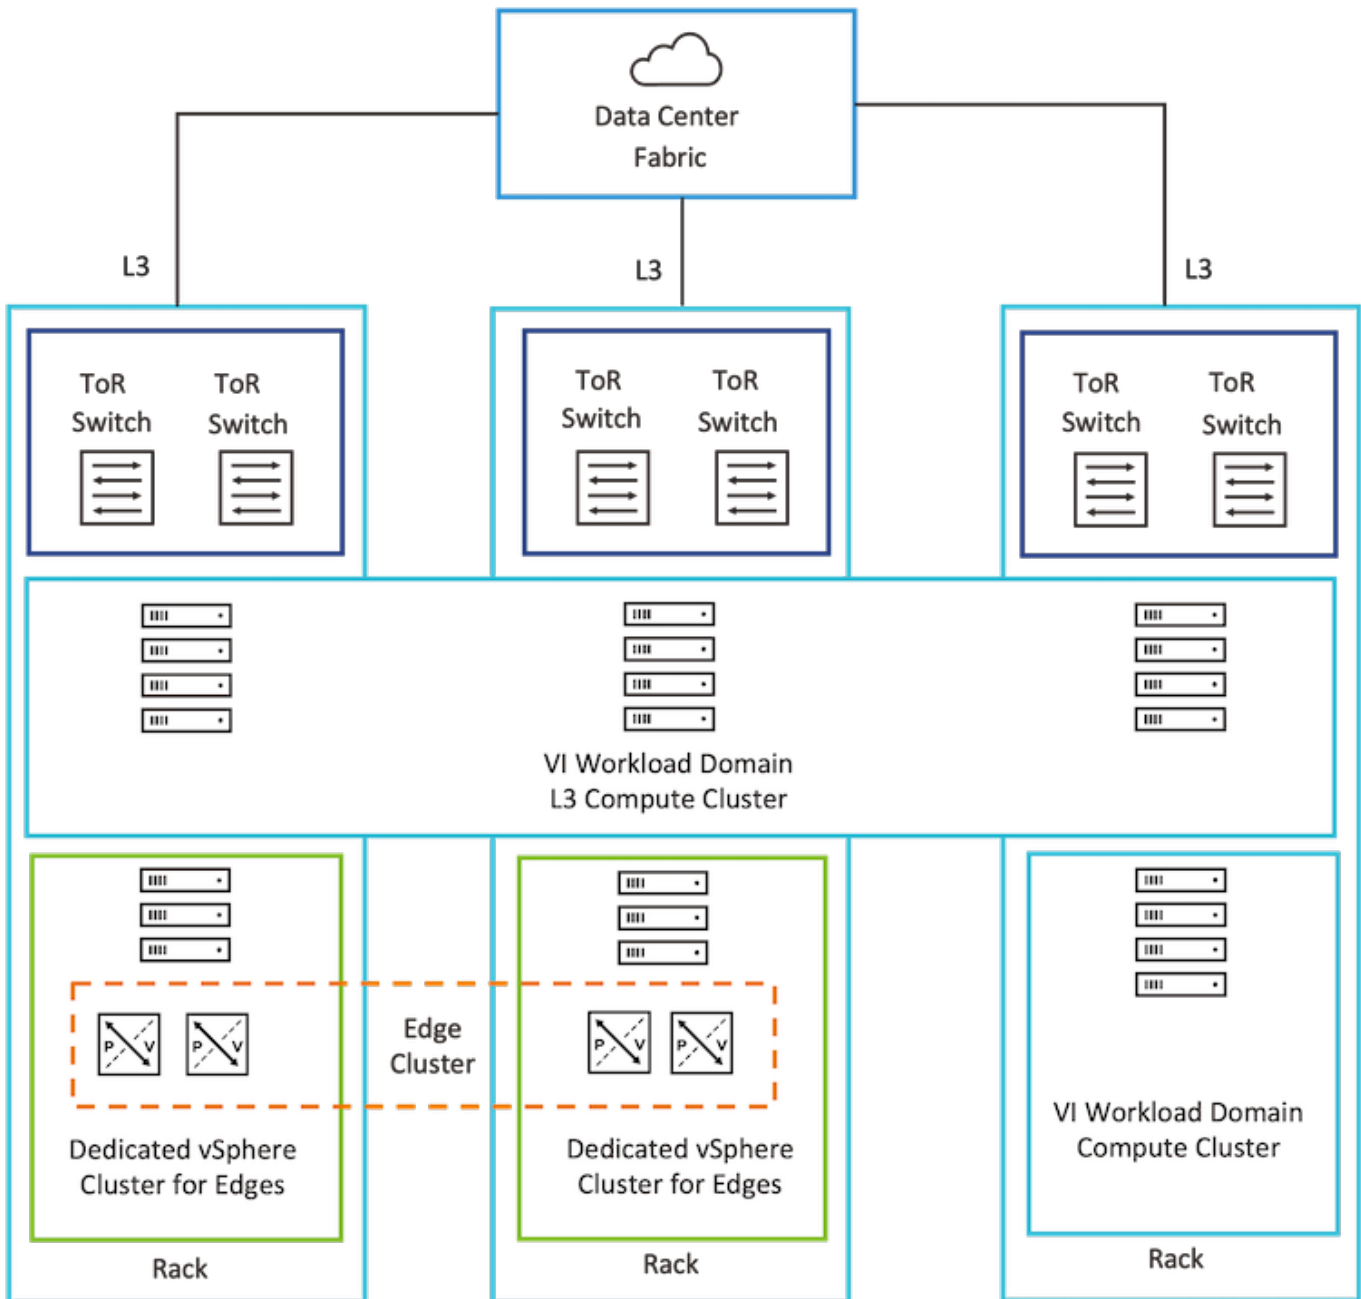
Table 77: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

NSX Edge Cluster Design Pattern Two: Multi-Rack Edge Availability

This design pattern lists the design choices and resulting requirements and recommendations to deploy an NSX Edge cluster with edges deployed to two dedicated vSphere clusters in two independent racks. This design provides rack resiliency and a high level of availability for the edge cluster with edge VMs in both racks in the VMware Cloud Foundation deployment.

Figure 21: Multi-Rack Edge Availability



Design Choices for NSX Edge Cluster Design Pattern Two

The following choices for its VMware Cloud Foundation deployment are listed here:

Table 78: Design Choices for NSX Edge Cluster Design Pattern Two

Design Aspect	Choice Made
Architecture model	Standard

Table continued on next page

Continued from previous page

Design Aspect	Choice Made
Workload domain type	VI workload domain
Topology	Single Instance - Single Availability Zone
Workload domain cluster to rack mapping	Workload domain cluster per rack
Physical network configuration	Leaf-spine
Workload domain principal storage	vSAN

Design Elements for NSX Edge Cluster Design Pattern Two

Table 79: External Services Design Elements

Design Area	Applicable Design Elements
External services	External Services Design Requirements

Table 80: Physical Network Design Elements

Design Area	Applicable Design Elements
Physical network	Leaf-Spine Physical Network Design Requirements
	Leaf-Spine Physical Network Design Recommendations

Table 81: Multi-Rack Edge Availability Design Elements

Design Area	Applicable Design Elements
vSAN	vSAN Design Requirements
	vSAN Design Recommendations
ESXi	ESXi Server Design Requirements
	ESXi Server Design Recommendations
vSphere cluster	vSphere Cluster Design Requirements
	vSphere Cluster Design Recommendations
vSphere networking	vSphere Networking Design Recommendations
NSX Edge Node	NSX Edge Design Requirements
	NSX Edge Design Requirements for Multi-Rack Edge Availability
	NSX Edge Design Recommendations
Routing	BGP Routing Design Requirements
	BGP Routing Design Requirements for Multi-Rack Edge Availability
	BGP Routing Design Recommendations
Overlay	Overlay Design Requirements
	Overlay Design Recommendations

Table 82: Life Cycle Management Design Elements

Design Area	Applicable Design Elements
Life cycle management	Life Cycle Management Design Requirements

Table 83: Account and Password Management Design Elements

Design Area	Applicable Design Elements
Account and password management	Account and Password Management Design Recommendations

Table 84: Certificate Management Design Elements

Design Area	Applicable Design Elements
Certificate management	Certificate Management Design Recommendations

Appendix: Design Elements for VMware Cloud Foundation

The appendix aggregates all design requirements and recommendations in the design guidance for VMware Cloud Foundation. You can use this list for reference related to the end state of your platform and potentially to track your level of adherence to the design and any justification for deviation.

Architecture Design Elements for VMware Cloud Foundation

Use this list of requirements for reference related to using the standard or consolidated architecture of VMware Cloud Foundation.

For full design details, see [Architecture Models and Workload Domain Types in VMware Cloud Foundation](#).

Table 85: Architecture Model Recommendations for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-ARCH-RCMD-CFG-001	Use the standard architecture model of VMware Cloud Foundation.	<ul style="list-style-type: none"> Aligns with the VMware best practice of separating management workloads from customer workloads. Provides better long-term flexibility and expansion options. 	Requires additional hardware.

Workload Domain Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the types of virtual infrastructure (VI) workload domains in a VMware Cloud Foundation environment.

For full design details, see [Workload Domain Types](#).

Table 86: Workload Domain Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WLD-RCMD-CFG-001	Use VI workload domains or isolated VI workload domains for customer workloads.	<ul style="list-style-type: none"> Aligns with the VMware best practice of separating management workloads from customer workloads. Provides better long term flexibility and expansion options. 	Requires additional hardware.

External Services Design Elements for VMware Cloud Foundation

Use this list of requirements for reference related to the configuration of external infrastructure services in an environment with a single or multiple VMware Cloud Foundation instances. The requirements define IP address allocation, name resolution, and time synchronization.

For full design details, see [External Services Design for VMware Cloud Foundation](#).

Table 87: External Services Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-EXT-REQD-NET-001	Allocate statically assigned IP addresses and host names for all workload domain components.	Ensures stability across the VMware Cloud Foundation instance, and makes it simpler to maintain, track, and implement a DNS configuration.	You must provide precise IP address management.
VCF-EXT-REQD-NET-002	Configure forward and reverse DNS records for all workload domain components.	Ensures that all components are accessible by using a fully qualified domain name instead of by using IP addresses only. It is easier to remember and connect to components across the VMware Cloud Foundation instance.	You must provide DNS records for each component.
VCF-EXT-REQD-NET-003	Configure time synchronization by using an internal NTP time source for all workload domain components.	Ensures that all components are synchronized with a valid time source.	An operational NTP service must be available in the environment.
VCF-EXT-REQD-NET-004	Set the NTP service for all workload domain components to start automatically.	Ensures that the NTP service remains synchronized after you restart a component.	None.

Physical Network Design Elements for VMware Cloud Foundation

Use this design decision list for reference related to the configuration of the physical network in an environment with a single or multiple VMware Cloud Foundation instances. The design also considers if an instance contains a single or multiple availability zones.

For full design details, see [Physical Network Infrastructure Design for](#) .

Table 88: Leaf-Spine Physical Network Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-REQD-CFG-001	Do not use EtherChannel (LAG, LACP, or vPC) configuration for ESXi host uplinks.	<ul style="list-style-type: none"> • Simplifies configuration of top-of-rack switches. • Teaming options available with vSphere Distributed Switch provide load balancing and failover. • EtherChannel implementations might have vendor-specific limitations. 	None.
VCF-NET-REQD-CFG-002	Use VLANs to separate physical network functions.	<ul style="list-style-type: none"> • Supports physical network connectivity without requiring many NICs. • Isolates the different network functions in the SDDC so that you can have differentiated services and prioritized traffic as needed. 	Requires uniform configuration and presentation on all the trunks that are made available to the ESXi hosts.
VCF-NET-REQD-CFG-003	Configure the VLANs as members of a 802.1Q trunk.	All VLANs become available on the same physical network adapters on the ESXi hosts.	Optionally, the management VLAN can act as the native VLAN.
VCF-NET-REQD-CFG-004	Set the MTU size to at least 1,700 bytes (recommended 9,000 bytes for jumbo frames) on the physical switch ports, vSphere Distributed Switches, vSphere Distributed Switch port groups, and N-VDS switches that support the following traffic types: <ul style="list-style-type: none"> • Overlay (Geneve) • vSAN • vSphere vMotion 	<ul style="list-style-type: none"> • Improves traffic throughput. • Supports Geneve by increasing the MTU size to a minimum of 1,600 bytes. • Geneve is an extensible protocol. The MTU size might increase with future capabilities. While 1,600 bytes is sufficient, an MTU size of 1,700 bytes provides more room for increasing the Geneve MTU size without the need to change the MTU 	<p>When adjusting the MTU packet size, you must also configure the entire network path (VMkernel network adapters, virtual switches, physical switches, and routers) to support the same MTU packet size.</p> <p>In an environment with multiple availability zones, the MTU must be configured on the entire network path between the zones.</p>

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		size of the physical infrastructure.	

Table 89: Leaf-Spine Physical Network Design Requirements for Multi-Rack Compute VI Workload Domain Cluster for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-L3MR-REQD-CFG-001	<p>For a multi-rack compute VI workload domain cluster, provide separate VLANs per rack for each network function.</p> <ul style="list-style-type: none"> • Host management • vSAN • vSphere vMotion • Host overlay 	A Layer 3 leaf-spine fabric has a Layer 3 boundary at the leaf switches in each rack creating a Layer 3 boundary between racks.	Requires additional VLANs for each rack.
VCF-NET-L3MR-REQD-CFG-002	<p>For a multi-rack compute VI workload domain cluster, the subnets for each network per rack must be routable and reachable to the leaf switches in the other racks.</p> <ul style="list-style-type: none"> • Host management • vSAN • vSphere vMotion • Host overlay 	Ensures the traffic for each network can flow between racks.	Requires additional physical network configuration to make networks routable between racks.

Table 90: Leaf-Spine Physical Network Design Requirements for Multi-Rack NSX Edge Availability for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-MRE-REQD-CFG-001	<p>For multi-rack NSX Edge availability, in each rack that is dedicated for edge nodes, configure the leaf switches with the following VLANs:</p> <ul style="list-style-type: none"> • VM management • Edge Uplink 1 • Edge Uplink 2 • Edge overlay 	A Layer 3 leaf-spine fabric has a Layer 3 boundary at the leaf switches in each rack creating a Layer 3 boundary between racks.	Requires additional VLANs for each rack.
VCF-NET-MRE-REQD-CFG-002	<p>For multi-rack NSX Edge availability, in each rack that is dedicated for edge nodes, the subnets for the following</p>	Ensures the traffic for the edge TEPs network can flow between racks.	Requires additional physical network configuration to ensure networks are routable between racks.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	<p>networks must be routable and reachable to the leaf switches in the other rack.</p> <ul style="list-style-type: none"> • VM management • Edge overlay 		

Table 91: Leaf-Spine Physical Network Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NET-REQD-CFG-005	<p>Set the MTU size to at least 1,500 bytes (1,700 bytes preferred; 9,000 bytes recommended for jumbo frames) on the components of the physical network between the VMware Cloud Foundation instances for the following traffic types.</p> <ul style="list-style-type: none"> • NSX Edge RTEP 	<ul style="list-style-type: none"> • Jumbo frames are not required between VMware Cloud Foundation instances. However, increased MTU improves traffic throughput. • Increasing the RTEP MTU to 1,700 bytes minimizes fragmentation for standard-size workload packets between VMware Cloud Foundation instances. 	When adjusting the MTU packet size, you must also configure the entire network path, that is, virtual interfaces, virtual switches, physical switches, and routers to support the same MTU packet size.
VCF-NET-REQD-CFG-006	<p>Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 500 ms.</p>	A latency lower than 500 ms is required for NSX Federation.	None.
VCF-NET-REQD-CFG-007	<p>Provide a routed connection between the NSX Manager clusters in VMware Cloud Foundation instances that are connected in an NSX Federation.</p>	Configuring NSX Federation requires connectivity between the NSX Global Manager instances, NSX Local Manager instances, and NSX Edge clusters.	You must assign unique routable IP addresses for each fault domain.

Table 92: Leaf-Spine Physical Network Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-001	<p>Use two ToR switches for each rack.</p>	Supports the use of two 10-GbE (25-GbE or greater recommended) links to each server, provides redundancy and reduces the overall design complexity.	Requires two ToR switches per rack which might increase costs.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-002	Implement the following physical network architecture: <ul style="list-style-type: none"> • At least one 25-GbE (10-GbE minimum) port on each ToR switch for ESXi host uplinks (Host to ToR). • Layer 3 device that supports BGP. 	<ul style="list-style-type: none"> • Provides availability during a switch failure. • Provides support for BGP dynamic routing protocol 	<ul style="list-style-type: none"> • Might limit the hardware choices. • Requires dynamic routing protocol configuration in the physical network.
VCF-NET-RCMD-CFG-003	Use a physical network that is configured for BGP routing adjacency.	<ul style="list-style-type: none"> • Supports design flexibility for routing multi-site and multi-tenancy workloads. • BGP is the only dynamic routing protocol that is supported for NSX Federation. • Supports failover between ECMP Edge uplinks. 	Requires BGP configuration in the physical network.
VCF-NET-RCMD-CFG-004	Assign persistent IP configurations for NSX tunnel endpoints (TEPs) that use static IP pools instead of dynamic IP pool addressing.	<ul style="list-style-type: none"> • Ensures that endpoints have a persistent TEP IP address. • In VMware Cloud Foundation, TEP IP assignment by using static IP pools is recommended for all topologies. • This configuration removes any requirement for external DHCP services. 	Adding more hosts to the cluster may require the static IP pools to be increased..
VCF-NET-RCMD-CFG-005	Configure the trunk ports connected to ESXi NICs as trunk PortFast.	Reduces the time to transition ports over to the forwarding state.	Although this design does not use the STP, switches usually have STP configured by default.
VCF-NET-RCMD-CFG-006	Configure VRRP, HSRP, or another Layer 3 gateway availability method for these networks. <ul style="list-style-type: none"> • Management • Edge overlay 	Ensures that the VLANs that are stretched between availability zones are connected to a highly-available gateway. Otherwise, a failure in the Layer 3 gateway will cause disruption in the traffic in the SDN setup.	Requires configuration of a high availability technology for the Layer 3 gateways in the data center.
VCF-NET-RCMD-CFG-007	Use separate VLANs for the network functions for each cluster.	Reduces the size of the Layer 2 broadcast domain to a single vSphere cluster.	Increases the overall number of VLANs that are

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
			required for a VMware Cloud Foundation instance.

Table 93: Leaf-Spine Physical Network Design Recommendations for Dedicated Edge Scale and Performance for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-DES-RCMD-CFG-001	Implement the following physical network architecture: <ul style="list-style-type: none"> • Two 100-GbE ports on each ToR switch for ESXi host uplinks (Host to ToR). • Layer 3 device that supports BGP. 	Supports the requirements for high bandwidth and packets per second for large-scale deployments.	Requires 100-GbE network switches.

Table 94: Leaf-Spine Physical Network Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NET-RCMD-CFG-008	Provide BGP routing between all VMware Cloud Foundation instances that are connected in an NSX Federation setup.	BGP is the supported routing protocol for NSX Federation.	None.
VCF-NET-RCMD-CFG-009	Ensure that the latency between VMware Cloud Foundation instances that are connected in an NSX Federation is less than 150 ms for workload mobility.	A latency lower than 150 ms is required for the following features: <ul style="list-style-type: none"> • Cross vCenter Server vMotion 	None.

vSAN Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to shared storage, vSAN principal storage, and NFS supplemental storage in an environment with a single or multiple VMware Cloud Foundation instances. The design also considers whether an instance contains a single or multiple availability zones.

After you set up the physical storage infrastructure, the configuration tasks for most design decisions are automated in VMware Cloud Foundation. You must perform the configuration manually only for a limited number of design elements as noted in the design implication.

For full design details, see [vSAN Design for VMware Cloud Foundation](#).

Table 95: vSAN Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-001	Provide sufficient raw capacity to meet the initial needs of the workload domain cluster.	Ensures that sufficient resources are present to create the workload domain cluster.	None.
VCF-VSAN-REQD-CFG-002	Provide at least the required minimum number of hosts according to the cluster type.	Satisfies the requirements for storage availability.	None.

Table 96: vSAN ESA Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-003	Verify the hardware components used in your vSAN deployment are on the vSAN Hardware Compatibility List.	Prevents hardware-related failures during workload deployment	Limits the number of compatible hardware configurations that can be used.

Table 97: vSAN Max Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-MAX-REQD-CFG-001	Provide at least four nodes for the initial cluster.	A vSAN Max cluster must contain at least four hosts.	None.

Table 98: vSAN Design Requirements for a Multi-Rack Compute VI Workload Domain Cluster with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-L3MR-REQD-CFG-001	Configure vSAN fault domains and place the nodes of each rack in their fault domain.	Allows workload VMs to tolerate a rack failure by distributing copies of the data and witness components on nodes in separate racks.	You must make the fault domain configuration manually vCenter Server after deployment and after cluster expansion.

Table 99: vSAN ESA Design Requirements for Multi-Rack Compute VI Workload Domain Cluster with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-ESA-L3MR-	Use a minimum of four racks for the cluster.	Provides support for reprotecting vSAN objects if a single-rack failure occurs.	Requires a minimum of four hosts in a cluster with RAID-5 erasure coding.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
REQD-CFG-001			
VCF-VSAN-ESA-L3MR-REQD-CFG-002	Deactivate vSAN ESA auto policy management.	<ul style="list-style-type: none"> Not supported with vSAN fault domains. Provides support for applying a compliant default storage policy on the vSAN datastore. 	To align with the number of vSAN fault domains, you might have to create a default storage policy manually.

Table 100: vSAN OSA Design Requirements for Multi-Rack Compute VI Workload Domain Cluster with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-OSA-L3MR-REQD-CFG-001	Use a minimum of four racks for the cluster.	Provides support for reprotecting vSAN objects if a single-rack failure occurs.	<ul style="list-style-type: none"> Requires a minimum of four hosts in a cluster with RAID-1 configuration.

Table 101: vSAN Design Requirements for Stretched Clusters with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VSAN-REQD-CFG-004	Add the following setting to the default vSAN storage policy: Site disaster tolerance = Site mirroring - stretched cluster	Provides the necessary protection for virtual machines in each availability zone, with the ability to recover from an availability zone outage.	You might need additional policies if third-party virtual machines are to be hosted in these clusters because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
VCF-VSAN-REQD-CFG-005	Configure two fault domains, one for each availability zone. Assign each host to their respective availability zone fault domain.	Fault domains are mapped to availability zones to provide logical host separation and ensure a copy of vSAN data is always available even when an availability zone goes offline.	You must provide additional raw storage when the site mirroring - stretched cluster option is selected, and fault domains are enabled.
VCF-VSAN-REQD-CFG-006	Configure an individual vSAN storage policy for each stretched cluster.	The vSAN storage policy of a stretched cluster cannot be shared with other clusters.	You must configure additional vSAN storage policies.
VCF-VSAN-WTN-REQD-CFG-001	Deploy a vSAN witness appliance in a location that is not local to the ESXi hosts in any of the availability zones.	Ensures availability of vSAN witness components in the event of a failure of one of the availability zones.	You must provide a third physically separate location that runs a vSphere environment. You might use a VMware Cloud Foundation instance in a separate physical location.
VCF-VSAN-WTN-REQD-CFG-002	Deploy a witness appliance that corresponds to the	Ensures the witness appliance is sized to support the projected workload storage consumption.	The vSphere environment at the witness location must satisfy the resource requirements of the witness appliance.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	required cluster capacity.		
VCF-VSAN-WTN-REQD-CFG-003	Connect the first VMkernel adapter of the vSAN witness appliance to the management network in the witness site.	Enables connecting the witness appliance to the workload domain vCenter Server.	The management networks in both availability zones must be routed to the management network in the witness site.
VCF-VSAN-WTN-REQD-CFG-004	Allocate a statically assigned IP address and host name to the management adapter of the vSAN witness appliance.	Simplifies maintenance and tracking, and implements a DNS configuration.	Requires precise IP address management.
VCF-VSAN-WTN-REQD-CFG-005	Configure forward and reverse DNS records for the vSAN witness appliance for the VMware Cloud Foundation instance.	Enables connecting the vSAN witness appliance to the workload domain vCenter Server by FQDN instead of IP address.	You must provide DNS records for the vSAN witness appliance.
VCF-VSAN-WTN-REQD-CFG-006	Configure time synchronization by using an internal NTP time for the vSAN witness appliance.	Prevents any failures in the stretched cluster configuration that are caused by time mismatch between the vSAN witness appliance and the ESXi hosts in both availability zones and workload domain vCenter Server.	<ul style="list-style-type: none"> An operational NTP service must be available in the environment. All firewalls between the vSAN witness appliance and the NTP servers must allow NTP traffic on the required network ports.

Table 102: vSAN Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-001	Provide sufficient raw capacity to meet the planned needs of the workload domain cluster.	Ensures that sufficient resources are present in the workload domain cluster, preventing the need to expand the vSAN datastore in the future.	None.
VCF-VSAN-RCMD-CFG-002	Ensure that at least 30% of free space is always available on the vSAN datastore.	This reserved capacity is set aside for host maintenance mode data evacuation, component rebuilds, rebalancing operations, and VM snapshots.	Increases the amount of available storage needed.
VCF-VSAN-RCMD-CFG-003	Use the default VMware vSAN storage policy.	<ul style="list-style-type: none"> Provides the level of redundancy that is needed in the workload domain cluster. Provides the level of performance that is enough for the individual workloads. 	You might need additional policies for third-party virtual machines hosted in these

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
			clusters because their performance or availability requirements might differ from what the default VMware vSAN policy supports.
VCF-VSAN-RCMD-CFG-004	Leave the default virtual machine swap file as a sparse object on vSAN.	Sparse virtual swap files consume capacity on vSAN only as they are accessed. As a result, you can reduce the consumption on the vSAN datastore if virtual machines do not experience memory over-commitment, which would require the use of the virtual swap file.	None.
VCF-VSAN-RCMD-CFG-005	Use the existing vSphere Distributed Switch instance for the workload domain cluster.	<ul style="list-style-type: none"> Reduces the complexity of the network design. Reduces the number of physical NICs required. 	All traffic types can be shared over common uplinks.
VCF-VSAN-RCMD-CFG-006	Configure jumbo frames on the VLAN for vSAN traffic.	<ul style="list-style-type: none"> Simplifies configuration because jumbo frames are also used to improve the performance of vSphere vMotion and NFS storage traffic. Reduces the CPU overhead, resulting in high network usage. 	Every device in the network must support jumbo frames.
VCF-VSAN-RCMD-CFG-007	Use a dedicated VLAN for vSAN traffic for each vSAN cluster.	<ul style="list-style-type: none"> Isolates the vSAN traffic only to the hosts in the cluster. Reduces the amount of vSAN broadcast traffic that reaches a host to only the hosts in the cluster. 	Increases the number of VLANs required.
VCF-VSAN-RCMD-CFG-008	Configure vSAN in an all-flash configuration in the default workload domain cluster.	Meets the performance needs of the default workload domain cluster.	All vSAN disks must be flash disks, which might cost more than magnetic disks.

Table 103: vSAN OSA Design Recommendations for with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-009	Ensure that the storage I/O controller has a minimum queue depth of 256 set.	Storage controllers with lower queue depths can cause performance and stability problems when running vSAN.	Limits the number of compatible I/O controllers

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		vSAN ReadyNode servers are configured with the correct queue depths for vSAN.	that can be used for storage.
VCF-VSAN-RCMD-CFG-010	Do not use the storage I/O controllers that are running vSAN disk groups for another purpose.	Running non-vSAN disks, for example, VMFS, on a storage I/O controller that is running a vSAN disk group can impact vSAN performance.	If non-vSAN disks are required in ESXi hosts, you must have an additional storage I/O controller in the host.
VCF-VSAN-RCMD-CFG-011	Configure vSAN with a minimum of two disk groups per ESXi host.	Reduces the size of the fault domain and spreads the I/O load over more disks for better performance.	Using multiple disk groups requires more disks in each ESXi host.
VCF-VSAN-RCMD-CFG-012	For the cache tier in each disk group, use a flash-based drive that is at least 600 GB large.	Provides enough cache for both hybrid or all-flash vSAN configurations to buffer I/O and ensure disk group performance. Additional space in the cache tier does not increase performance.	Using larger flash disks can increase the initial host cost.

Table 104: vSAN ESA Design Recommendations for with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-RCMD-CFG-013	Activate auto-policy management.	Configures optimized storage policies based on the cluster type and the number of hosts in the cluster inventory. Changes to the number of hosts in the cluster or Host Rebuild Reserve will prompt you to make a suggested adjustment to the optimized storage policy.	None.
VCF-VSAN-RCMD-CFG-014	Activate vSAN ESA compression.	Improves performance.	PostgreSQL databases and other applications might use their own compression capabilities. In these cases, using a storage policy with the

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
			compression capability turned off will save CPU cycles. You can disable vSAN ESA compressions for such workloads by using the Storage Policy Based Management (SPBM) framework.
VCF-VSAN-RCMD-CFG-014	Use NICs with a minimum 25-GbE capacity.	10-GbE NICs will limit the scale and performance of a vSAN ESA cluster because usually performance requirements increase over the lifespan of the cluster.	Requires 25-GbE or faster network fabric.

Table 105: vSAN Max Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-MAX-RCMD-001	Limit the size of the vSAN Max cluster to 24 hosts.	The total number of hosts for the vSAN Max cluster and vSAN compute clusters mounting the datastore must not exceed 128 hosts. A vSAN Max cluster with 24 nodes provides support for up to 104 vSAN compute hosts which is a good compute-to-storage ratio.	Limits the maximum number of hosts.
VCF-VSAN-MAX-RCMD-002	If the vSAN Max cluster consists of only four hosts, do not enable the Host Rebuild Reserve.	If the auto-policy management feature is turned on, prevents vSAN Max from using adaptive RAID-5.	vSAN does not reserve any of this capacity and presents it as free capacity for vSAN to self-repair if a single host failure occurs.
VCF-VSAN-MAX-RCMD-003	Include TPM (Trusted Platform Module) in the hardware configuration of the hosts.	Ensures that the keys issued to the hosts in a vSAN Max cluster using Data-at-Rest Encryption are cryptographically stored on a TPM device.	Limits the choice of hardware configurations.

Table 106: vSAN Design Recommendations for Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VSAN-WTN-RCMD-CFG-001	Configure the vSAN witness appliance to use the first VMkernel adapter, that is the management interface, for vSAN witness traffic.	Removes the requirement to have static routes on the witness appliance as witness traffic is routed over the management network.	The management networks in both availability zones must be routed to the management network in the witness site.
VCF-VSAN-WTN-RCMD-CFG-002	Place witness traffic on the management VMkernel adapter of all the ESXi hosts in the workload domain.	Separates the witness traffic from the vSAN data traffic. Witness traffic separation provides the following benefits: <ul style="list-style-type: none"> Removes the requirement to have static routes from the vSAN networks in both availability zones to the witness site. Removes the requirement to have jumbo frames enabled on the path between each availability zone and the witness site because witness traffic can use a regular MTU size of 1500 bytes. 	The management networks in both availability zones must be routed to the management network in the witness site.

ESXi Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the ESXi host configuration in an environment with a single or multiple VMware Cloud Foundation instances. The design elements determine the ESXi hardware configuration, networking, life cycle management and remote access.

The configuration tasks for most design requirements and recommendations are automated in VMware Cloud Foundation. You must perform the configuration manually only for a limited number of decisions as noted in the design implications.

For full design details, see [ESXi Design for](#) .

Table 107: Design Requirements for ESXi Server Hardware

Requirement ID	Design Requirement	Requirement Justification	Requirement Implication
VCF-ESX-REQD-CFG-001	Install no less than the minimum number of ESXi hosts required for the cluster type being deployed.	<ul style="list-style-type: none"> Ensures availability requirements are met. If one of the hosts is not available because of a failure or maintenance event, the CPU overcommitment ratio becomes 2:1. 	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Requirement Justification	Requirement Implication
VCF-ESX-REQD-CFG-002	Ensure each ESXi host matches the required CPU, memory and storage specification.	<ul style="list-style-type: none"> Ensures workloads will run without contention even during failure and maintenance conditions. 	Assemble the server specification and number according to the sizing in VMware Cloud Foundation Planning and Preparation Workbook which is based on projected deployment size.
VCF-ESX-REQD-SEC-001	Regenerate the certificate of each ESXi host after assigning the host an FQDN.	Establishes a secure connection with VMware Cloud Builder during the deployment of a workload domain and prevents man-in-the-middle (MITM) attacks.	You must manually regenerate the certificates of the ESXi hosts before the deployment of a workload domain.

Table 108: Design Recommendations for ESXi Server Hardware

Recommendation ID	Recommendation	Justification	Implication
VCF-ESX-RCMD-CFG-001	Use vSAN ReadyNodes with vSAN storage for each ESXi host in the management domain.	<p>Your management domain is fully compatible with vSAN at deployment.</p> <p>For information about the models of physical servers that are vSAN-ready, see vSAN Compatibility Guide for vSAN ReadyNodes.</p>	<p>Hardware choices might be limited.</p> <p>If you plan to use a server configuration that is not a vSAN ReadyNode, your CPU, disks and I/O modules must be listed on the VMware Compatibility Guide under <i>CPU Series</i> and <i>vSAN Compatibility List</i> aligned to the ESXi version specified in VMware Cloud Foundation 5.2 Release Notes.</p>
VCF-ESX-RCMD-CFG-002	Allocate hosts with uniform configuration across the default management vSphere cluster.	<p>A balanced cluster has these advantages:</p> <ul style="list-style-type: none"> Predictable performance even during hardware failures Minimal impact of resynchronization or rebuild operations on performance 	You must apply vendor sourcing, budgeting, and procurement considerations for uniform server nodes on a per cluster basis.
VCF-ESX-RCMD-CFG-003	When sizing CPU, do not consider multithreading technology and associated performance gains.	Although multithreading technologies increase CPU performance, the performance gain depends on running workloads and	Because you must provide more physical CPU cores, costs increase and hardware choices become limited.

Table continued on next page

Continued from previous page

Recommendation ID	Recommendation	Justification	Implication
		differs from one case to another.	
VCF-ESX-RCMD-CFG-004	Install and configure all ESXi hosts in the default management cluster to boot using a 128-GB device or larger.	Provides hosts that have large memory, that is, greater than 512 GB, with enough space for the scratch partition when using vSAN.	None.
VCF-ESX-RCMD-CFG-005	Use the default configuration for the scratch partition on all ESXi hosts in the default management cluster.	<ul style="list-style-type: none"> If a failure in the vSAN cluster occurs, the ESXi hosts remain responsive and log information is still accessible. It is not possible to use vSAN datastore for the scratch partition. 	None.
VCF-ESX-RCMD-CFG-006	For workloads running in the default management cluster, save the virtual machine swap file at the default location.	Simplifies the configuration process.	Increases the amount of replication traffic for management workloads that are recovered as part of the disaster recovery process.
VCF-ESX-RCMD-NET-001	Place the ESXi hosts in each management domain cluster on a host management network that is separate from the VM management network.	<p>Enables the separation of the physical VLAN between ESXi hosts and the other management components for security reasons.</p> <p>The VM management network is not required for a multi-rack compute-only cluster in a VI workload domain.</p>	Increases the number of VLANs required.
VCF-ESX-RCMD-NET-002	Place the ESXi hosts in each VI workload domain on a separate host management VLAN-backed network.	Enables the separation of the physical VLAN between the ESXi hosts in different VI workload domains for security reasons.	Increases the number of VLANs required. For each VI workload domain, you must allocate a separate management subnet.
VCF-ESX-RCMD-SEC-001	Deactivate SSH access on all ESXi hosts in the management domain by having the SSH service stopped and using the default SSH service policy <code>Start</code> and <code>stop</code> manually.	<p>Ensures compliance with the <i>vSphere Security Configuration Guide</i> and with security best practices.</p> <p>Disabling SSH access reduces the risk of security attacks on the ESXi hosts through the SSH interface.</p>	You must activate SSH access manually for troubleshooting or support activities as VMware Cloud Foundation deactivates SSH on ESXi hosts after workload domain deployment.
VCF-ESX-RCMD-SEC-002	Set the advanced setting <code>UserVars.SuppressShell</code>	<ul style="list-style-type: none"> Ensures compliance with the <i>vSphere Security</i> 	You must turn off SSH enablement warning messages manually when

Table continued on next page

Continued from previous page

Recommendation ID	Recommendation	Justification	Implication
	Warning to 0 across all ESXi hosts in the management domain.	<p><i>Configuration Guide</i> and with security best practices</p> <ul style="list-style-type: none"> Enables the warning message that appears in the vSphere Client every time SSH access is activated on an ESXi host. 	performing troubleshooting or support activities.

vCenter Server Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the vCenter Server configuration in an environment with a single or multiple VMware Cloud Foundation instances. The design elements also consider if an instance contains a single or multiple availability zones. The vCenter Server design also includes the configuration of the default management cluster.

The configuration tasks for most design requirements and recommendations are automated in VMware Cloud Foundation. You must perform the configuration manually only for a limited number of decisions as noted in the design implications.

For full design details, see [vCenter Server Design for](#) .

vCenter Server Design Elements

Table 109: vCenter Server Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-CFG-001	Deploy a dedicated vCenter Server appliance for the management domain of the VMware Cloud Foundation instance.	<ul style="list-style-type: none"> Isolates vCenter Server failures to management or customer workloads. Isolates vCenter Server operations between management and customers. Supports a scalable cluster design where you can reuse the management components as more customer workloads are added to the SDDC. Simplifies capacity planning for customer workloads because you do not consider management workloads for the VI workload domain vCenter Server. Improves the ability to upgrade the vSphere environment and related components by enabling for explicit separation of maintenance windows: 	Requires a separate license for the vCenter Server instance in the management domain

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		<ul style="list-style-type: none"> – Management workloads remain available while you are upgrading the tenant workloads – Customer workloads remain available while you are upgrading the management nodes • Supports clear separation of roles and responsibilities to ensure that only administrators with granted authorization can control the management workloads. • Facilitates quicker troubleshooting and problem resolution. • Simplifies disaster recovery operations by supporting a clear separation between recovery of the management components and tenant workloads. • Provides isolation of potential network issues by introducing network separation of the clusters in the SDDC. 	
VCF-VCS-REQD-NET-001	Place all workload domain vCenters Server appliances on the VM management network in the management domain.	<ul style="list-style-type: none"> • Simplifies IP addressing for management VMs by using the same VLAN and subnet. • Provides simplified secure access to management VMs in the same VLAN network. 	None.

Table 110: vCenter Server Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-001	Deploy an appropriately sized vCenter Server appliance for each workload domain.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Small and for VI workload domains is Medium. To override these values, you must use the Cloud Builder API and the SDDC Manager API.
VCF-VCS-RCMD-CFG-002	Deploy a vCenter Server appliance with the appropriate storage size.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Small and for VI Workload Domains is Medium . To override these values, you must use the API.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-003	Protect workload domain vCenter Server appliances by using vSphere HA.	vSphere HA is the only supported method to protect vCenter Server availability in VMware Cloud Foundation.	vCenter Server becomes unavailable during a vSphere HA failover.
VCF-VCS-RCMD-CFG-004	In vSphere HA, set the restart priority policy for the vCenter Server appliance to high.	vCenter Server is the management and control plane for physical and virtual infrastructure. In a vSphere HA event, to ensure the rest of the SDDC management stack comes up faultlessly, the workload domain vCenter Server must be available first, before the other management components come online.	If the restart priority for another virtual machine is set to highest, the connectivity delay for the management components will be longer.

Table 111: vCenter Server Design Recommendations for vSAN Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VCS-RCMD-CFG-005	Add the vCenter Server appliance to the virtual machine group for the first availability zone.	Ensures that, by default, the vCenter Server appliance is powered on a host in the first availability zone.	None.

vCenter Single Sign-On Design Elements

Table 112: Design Requirements for the Multiple vCenter Server Instance - Single vCenter Single Sign-on Domain Topology for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-SSO-STD-001	Join all vCenter Server instances within a VMware Cloud Foundation instance to a single vCenter Single Sign-On domain.	When all vCenter Server instances are in the same vCenter Single Sign-On domain, they can share authentication and license data across all components.	<ul style="list-style-type: none"> Only one vCenter Single Sign-On domain exists. The number of linked vCenter Server instances in the same vCenter Single Sign-On domain is limited to 15 instances. Because each workload domain uses a dedicated vCenter Server instance, you can deploy up to 15 domains within each VMware Cloud Foundation instance.
VCF-VCS-REQD-SSO-STD-002	Create a ring topology between the vCenter Server	By default, one vCenter Server instance replicates only with another vCenter	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	instances within the VMware Cloud Foundation instance.	Server instance. This setup creates a single point of failure for replication. A ring topology ensures that each vCenter Server instance has two replication partners and removes any single point of failure.	

Table 113: Design Requirements for Multiple vCenter Server Instance - Multiple vCenter Single Sign-On Domain Topology for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VCS-REQD-SSO-ISO-001	Create all vCenter Server instances within a VMware Cloud Foundation instance in their own unique vCenter Single Sign-On domains.	<ul style="list-style-type: none"> Enables isolation at the vCenter Single Sign-On domain layer for increased security separation. Supports up to 25 workload domains. 	<ul style="list-style-type: none"> Each vCenter server instance is managed through its own pane of glass using a different set of administrative credentials. You must manage password rotation for each vCenter Single Sign-On domain separately.

vSphere Cluster Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the vSphere cluster configuration in an environment with a single or multiple VMware Cloud Foundation instances. The design elements also consider if an instance contains a single or multiple availability zones.

For full design details, see [Logical vSphere Cluster Design for](#) .

Table 114: vSphere Cluster Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-001	Create a cluster in each workload domain for the initial set of ESXi hosts.	<ul style="list-style-type: none"> Simplifies configuration by isolating management from customer workloads. Ensures that customer workloads have no impact on the management stack. 	Management of multiple clusters and vCenter Server instances increases operational overhead.
VCF-CLS-REQD-CFG-002	Allocate a minimum number of ESXi hosts according to the cluster type being deployed.	<ul style="list-style-type: none"> Ensures correct level of redundancy to protect against host failure in the cluster. 	To support redundancy, you must allocate additional ESXi host resources.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-003	<p>If using a consolidated workload domain, configure the following vSphere resource pools to control resource usage by management and customer workloads.</p> <ul style="list-style-type: none"> • <code>cluster-name-rp-sddc-mgmt</code> • <code>cluster-name-rp-sddc-edge</code> • <code>cluster-name-rp-user-edge</code> • <code>cluster-name-rp-user-vm</code> 	<ul style="list-style-type: none"> • Ensures sufficient resources for the management components. 	You must manage the vSphere resource pool settings over time.
VCF-CLS-REQD-CFG-004	For vSAN clusters, except for vSAN Max clusters, configure the vSAN network gateway IP address as the isolation address for the cluster.	vSphere HA can validate if a host is isolated from the vSAN network.	You must allocate an additional IP address.
VCF-CLS-REQD-CFG-005	For vSAN clusters, except for vSAN Max clusters, set the advanced cluster setting <code>das.usedefaultisolationaddress</code> to false.	Ensures that vSphere HA uses the manual isolation addresses instead of the default management network gateway address.	None.

Table 115: vSphere Cluster Design Requirements for vSAN Stretched Clusters with VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-REQD-CFG-006	Configure the IP address of the vSAN network for the second availability zone as an additional isolation addresses for the cluster.	Enables vSphere HA to validate if a host is isolated from the vSAN network for hosts in both availability zones.	The IP address of the vSAN network gateway must be highly available and reply to ICMP requests.
VCF-CLS-REQD-CFG-007	Enable the Override default gateway for this adapter setting on the vSAN VMkernel adapters on all ESXi hosts.	Enables routing the vSAN data traffic through the vSAN network gateway rather than through the management gateway.	vSAN networks across availability zones must have a route to each other.
VCF-CLS-REQD-CFG-008	Create a host group for each availability zone and add the ESXi hosts in the zone to the respective group.	Makes it easier to manage which virtual machines run in which availability zone.	You must create and maintain VM-Host DRS group rules.

Table 116: vSphere Cluster Design Requirements for a Multi-Rack Compute VI Workload Domain Cluster for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-L3MR-REQD-CFG-001	Configure the IP address of the vSAN network gateway for each rack accessible over Layer 3 as the isolation address for the nodes in that rack in the cluster.	Enables vSphere HA to validate if a host is isolated from the vSAN network.	The IP address of the vSAN network gateway must be highly available and reply to ICMP requests.

Table 117: vSphere Cluster Design Requirements for Multi-Rack Edge Availability for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-MRE-REQD-ENV-001	Deploy a minimum of two vSphere clusters for NSX edge nodes, with one vSphere cluster in each rack.	Provides availability if a failure of a rack or single vSphere cluster occurs.	Additional cluster required.

Table 118: vSphere Cluster Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-001	Use vSphere HA to protect all virtual machines against failures.	vSphere HA supports a robust level of protection for both ESXi host and virtual machine availability.	You must provide sufficient resources on the remaining hosts so that virtual machines can be restarted on those hosts in the event of a host outage.
VCF-CLS-RCMD-CFG-002	For vSAN clusters, set host isolation response to Power Off and restart VMs in vSphere HA.	vSAN requires that the host isolation response be set to Power Off and to restart virtual machines on available ESXi hosts.	If a false positive event occurs, virtual machines are powered off and an ESXi host is declared isolated incorrectly.
VCF-CLS-RCMD-CFG-003	Configure admission control for 1 ESXi host failure and percentage-based failover capacity.	Using the percentage-based reservation works well in situations where virtual machines have varying and sometimes significant CPU or memory reservations. vSphere automatically calculates the reserved percentage according to the number of ESXi host failures to tolerate and the number of ESXi hosts in the cluster.	In a cluster of 4 ESXi hosts, the resources of only 3 ESXi hosts are available for use.
VCF-CLS-RCMD-CFG-004	Enable VM Monitoring for each cluster.	VM Monitoring provides in-guest protection for most VM workloads. The application or service running on the	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		virtual machine must be capable of restarting successfully after a reboot or the virtual machine restart is not sufficient.	
VCF-CLS-RCMD-CFG-005	Set the advanced cluster setting <code>das.iostatsinterval</code> to 0 to deactivate monitoring the storage and network I/O activities of the management appliances.	Enables triggering a restart of a management appliance when an OS failure occurs and heartbeats are not received from VMware Tools instead of waiting additionally for the I/O check to complete.	If you want to specifically enable I/O monitoring, you must configure the das.iostatsinterval advanced setting.
VCF-CLS-RCMD-CFG-006	Enable vSphere DRS on all clusters, using the default fully automated mode with medium threshold.	Provides the best trade-off between load balancing and unnecessary migrations with vSphere vMotion.	If a vCenter Server outage occurs, the mapping from virtual machines to ESXi hosts might be difficult to determine.
VCF-CLS-RCMD-CFG-007	Enable Enhanced vMotion Compatibility (EVC) on all clusters in the management domain.	Supports cluster upgrades without virtual machine downtime.	You must enable EVC only if the clusters contain hosts with CPUs from the same vendor. You must enable EVC on the default management domain cluster during bringup.
VCF-CLS-RCMD-CFG-008	Set the cluster EVC mode to the highest available baseline that is supported for the lowest CPU architecture on the hosts in the cluster.	Supports cluster upgrades without virtual machine downtime.	None.
VCF-CLS-RCMD-LCM-001	Use images as the life cycle management method for all workload domains.	<ul style="list-style-type: none"> vSphere Lifecycle Manager images simplify the management of firmware and vendor additions manually. Supports vSAN ESA clusters. 	<ul style="list-style-type: none"> A cluster image is required during VI workload domain or cluster deployment. A cluster image is required when you add a cluster to the management domain.

Table 119: vSphere Cluster Design Recommendations for vSAN Stretched Clusters with VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-CLS-RCMD-CFG-009	Increase admission control percentage to half of the ESXi hosts in the cluster.	Allocating only half of a stretched cluster ensures that all VMs have enough	In a cluster of 8 ESXi hosts, the resources of only 4 ESXi hosts are available for use.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		resources if an availability zone outage occurs.	If you add more ESXi hosts to the default management cluster, add them in pairs, one per availability zone.
VCF-CLS-RCMD-CFG-010	Create a virtual machine group for each availability zone and add the VMs in the zone to the respective group.	Ensures that virtual machines are located only in the assigned availability zone to avoid unnecessary vSphere vMotion migrations.	You must add virtual machines to the allocated group manually.
VCF-CLS-RCMD-CFG-011	Create a should-run-on-hosts-in-group VM-Host affinity rule to run each group of virtual machines on the respective group of hosts in the same availability zone.	Ensures that virtual machines are located only in the assigned availability zone to avoid unnecessary vSphere vMotion migrations.	You must manually create the rules.

Table 120: vSphere Cluster Design Recommendations for a Multi-Rack Compute VI Workload Domain Cluster for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-CLS-L3MR-RCMD-CFG-001	Deploy a multi-rack VI workload domain cluster with Layer 3 networking in a minimum of four racks with a Layer 3 boundary at the rack level.	Improves resiliency. When combined with vSAN fault domains, the minimum of four racks protects against a single rack failure.	<ul style="list-style-type: none"> Requires a minimum of four hosts in the cluster. vSAN fault domains must be manually configured.

vSphere Networking Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the configuration of the vSphere Distributed Switch instances and VMkernel adapters in a VMware Cloud Foundation environment.

For full design details, see [vSphere Networking Design for](#) .

Table 121: vSphere Networking Design Requirements for a Multi-Rack Compute VI Workload Domain Cluster for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VDS-L3MR-REQD-CFG-001	<p>For each rack, create a port group on the vSphere Distributed Switch for the cluster for the following traffic types:</p> <ul style="list-style-type: none"> Host management vSAN vSphere vMotion 	Required for using separate VLANs per rack.	None.

Table 122: vSphere Networking Design Requirements for Dedicated Edge Scale and Performance for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VDS-DES-REQD-CFG-001	Configure Enhanced Datapath - Interrupt Mode	Provides the best performance for bandwidth and packets per second for the edge nodes running on the cluster.	The physical NIC must support the Enhanced Datapath - Interrupt Mode feature.
VCF-VDS-DES-REQD-CFG-002	Deactivate Network I/O Control on a vSphere Distributed Switch used for edge VM traffic.	Maximizes the packet per second rate that the edge nodes can achieve.	You must deactivate Network I/O Control manually after the cluster is deployed.

Table 123: vSphere Networking Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VDS-RCMD-CFG-001	Use a single vSphere Distributed Switch per cluster.	Reduces the complexity of the network design.	Increases the number of vSphere Distributed Switches that must be managed.
VCF-VDS-RCMD-CFG-002	Do not share a vSphere Distributed Switch across clusters.	<ul style="list-style-type: none"> Enables independent lifecycle management of vSphere Distributed Switch per cluster. Reduces the size of the fault domain. 	For multiple clusters, you manage more vSphere Distributed Switches .
VCF-VDS-RCMD-CFG-003	Configure the MTU size of the vSphere Distributed Switch to 9000 bytes for jumbo frames.	<ul style="list-style-type: none"> Supports the MTU size required by system traffic types. Improves traffic throughput. 	When adjusting the MTU packet size, you must also configure the entire network path (VMkernel ports, virtual switches, physical switches, and routers) to support the same MTU packet size.
VCF-VDS-RCMD-DPG-001	Use ephemeral port binding for the VM management port group.	<p>Using ephemeral port binding provides the option for recovery of the vCenter Server instance that is managing the distributed switch.</p> <p>The VM management network is not required for a multi-rack compute-only cluster in a VI workload domain.</p>	Port-level permissions and controls are lost across power cycles, and no historical context is saved.
VCF-VDS-RCMD-DPG-002	Use static port binding for all non-management port groups.	Static binding ensures a virtual machine connects to the same port on the vSphere Distributed Switch. This allows for historical	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		data and port level monitoring.	
VCF-VDS-RCMD-DPG-003	Use the Route based on physical NIC load teaming algorithm for the VM management port group. The VM management port group is not required for a multi-rack compute VI workload domain cluster.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-004	Use the Route based on physical NIC load teaming algorithm for the ESXi management port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-005	Use the Route based on physical NIC load teaming algorithm for the vSphere vMotion port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-DPG-006	Use the Route based on physical NIC load teaming algorithm for the vSAN port group.	Reduces the complexity of the network design, increases resiliency, and can adjust to fluctuating workloads.	None.
VCF-VDS-RCMD-NIO-001	Enable Network I/O Control on vSphere Distributed Switch of the management domain cluster.	Increases resiliency and performance of the network.	Network I/O Control might impact network performance for critical traffic types if misconfigured.
VCF-VDS-RCMD-NIO-002	Set the share value for management traffic to Normal.	By keeping the default setting of Normal, management traffic is prioritized higher than vSphere vMotion but lower than vSAN traffic. Management traffic is important because it ensures that the hosts can still be managed during times of network contention.	None.
VCF-VDS-RCMD-NIO-003	Set the share value for vSphere vMotion traffic to Low.	During times of network contention, vSphere vMotion traffic is not as important as virtual machine or storage traffic.	During times of network contention, vMotion takes longer than usual to complete.
VCF-VDS-RCMD-NIO-004	Set the share value for virtual machines to High.	Virtual machines are the most important asset in the SDDC. Leaving the default	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		setting of High ensures that they always have access to the network resources they need.	
VCF-VDS-RCMD-NIO-005	Set the share value for vSAN traffic to High.	During times of network contention, vSAN traffic needs guaranteed bandwidth to support virtual machine performance.	None.
VCF-VDS-RCMD-NIO-006	Set the share value for other traffic types to Low.	By default, V VMware Cloud Foundation does not use other traffic types, like vSphere FT traffic. Hence, these traffic types can be set the lowest priority.	None.

NSX Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the configuration of NSX in an environment with a single or multiple VMware Cloud Foundation instances. The design also considers if an instance contains a single or multiple availability zones.

For full design details, see [NSX Design for VMware Cloud Foundation](#).

NSX Manager Design Elements

Table 124: NSX Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LM-REQD-CFG-001	Place the appliances of the NSX Manager cluster on the VM management network in the management domain.	<ul style="list-style-type: none"> Simplifies IP addressing for management VMs by using the same VLAN and subnet. Provides simplified secure access to management VMs in the same VLAN network. 	None.
VCF-NSX-LM-REQD-CFG-002	Deploy three NSX Manager nodes in the default vSphere cluster in the management domain for configuring and managing the network services for the workload domain.	Supports high availability of the NSX manager cluster.	You must have sufficient resources in the default cluster of the management domain to run three NSX Manager nodes.

Table 125: NSX Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-001	Deploy appropriately sized nodes in the NSX Manager cluster for the workload domain.	Ensures resource availability and usage efficiency per workload domain.	The default size for a management domain is Medium, and for VI workload domains is Large.
VCF-NSX-LM-RCMD-CFG-002	Create a virtual IP (VIP) address for the NSX Manager cluster for the workload domain.	Provides high availability of the user interface and API of NSX Manager.	<ul style="list-style-type: none"> The VIP address feature provides high availability only. It does not load-balance requests across the cluster. When using the VIP address feature, all NSX Manager nodes must be deployed on the same Layer 2 network.
VCF-NSX-LM-RCMD-CFG-003	Apply VM-VM anti-affinity rules in vSphere Distributed Resource Scheduler (vSphere DRS) to the NSX Manager appliances.	Keeps the NSX Manager appliances running on different ESXi hosts for high availability.	You must allocate at least four physical hosts so that the three NSX Manager appliances continue running if an ESXi host failure occurs.
VCF-NSX-LM-RCMD-CFG-004	In vSphere HA, set the restart priority policy for each NSX Manager appliance to high.	<ul style="list-style-type: none"> NSX Manager implements the control plane for virtual network segments. vSphere HA restarts the NSX Manager appliances first so that other virtual machines that are being powered on or migrated by using vSphere vMotion while the control plane is offline lose connectivity only until the control plane quorum is re-established. Setting the restart priority to high reserves the highest priority for flexibility for adding services that must be started before NSX Manager. 	If the restart priority for another management appliance is set to highest, the connectivity delay for management appliances will be longer.

Table 126: NSX Manager Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-LM-RCMD-CFG-006	Add the NSX Manager appliances to the virtual	Ensures that, by default, the NSX Manager appliances	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
	machine group for the first availability zone.	are powered on a host in the primary availability zone.	

NSX Global Manager Design Elements

You must perform configuration tasks manually for the design requirements and recommendations for NSX Global Manager.

Table 127: NSX Global Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-GM-REQD-CFG-001	Place the appliances of the NSX Global Manager cluster on the Management VM network in each VMware Cloud Foundation instance.	<ul style="list-style-type: none"> Simplifies IP addressing for management VMs. Provides simplified secure access to all management VMs in the same VLAN network. 	None.

Table 128: NSX Global Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-001	Deploy three NSX Global Manager nodes for the workload domain to support NSX Federation across VMware Cloud Foundation instances.	Provides high availability for the NSX Global Manager cluster.	You must have sufficient resources in the default cluster of the management domain to run three NSX Global Manager nodes.
VCF-NSX-GM-RCMD-CFG-002	Deploy appropriately sized nodes in the NSX Global Manager cluster for the workload domain. Check VMware Configuration Maximums to select the right NSX Global Managers form factor for your scale needs.	Ensures resource availability and usage efficiency per workload domain.	Incorrectly sized NSX Global Managers might impact the ability to scale according to the requirements.
VCF-NSX-GM-RCMD-CFG-003	Create a virtual IP (VIP) address for the NSX Global Manager cluster for the workload domain.	Provides high availability of the user interface and API of NSX Global Manager.	<ul style="list-style-type: none"> The VIP address feature provides high availability only. It does not load-balance requests across the cluster. When using the VIP address feature, all NSX Global Manager nodes must be deployed on the same Layer 2 network.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-004	Apply VM-VM anti-affinity rules in vSphere DRS to the NSX Global Manager appliances.	Keeps the NSX Global Manager appliances running on different ESXi hosts for high availability.	You must allocate at least four physical hosts so that the three NSX Manager appliances continue running if an ESXi host failure occurs.
VCF-NSX-GM-RCMD-CFG-005	In vSphere HA, set the restart priority policy for each NSX Global Manager appliance to medium.	<ul style="list-style-type: none"> NSX Global Manager implements the management plane for global segments and firewalls. <p>NSX Global Manager is not required for control plane and data plane connectivity.</p> <ul style="list-style-type: none"> Setting the restart priority to medium reserves the high priority for services that impact the NSX control or data planes. 	<ul style="list-style-type: none"> Management of NSX global components will be unavailable until the NSX Global Manager virtual machines restart. The NSX Global Manager cluster is deployed in the management domain, where the total number of virtual machines is limited and where it competes with other management components for restart priority.
VCF-NSX-GM-RCMD-CFG-006	Deploy an additional NSX Global Manager Cluster in the second VMware Cloud Foundation instance.	Enables recoverability of NSX Global Manager in the second VMware Cloud Foundation instance if a failure in the first VMware Cloud Foundation instance occurs.	Requires additional NSX Global Manager nodes in the second VMware Cloud Foundation instance.
VCF-NSX-GM-RCMD-CFG-007	Set the NSX Global Manager cluster in the second VMware Cloud Foundation instance as standby for the workload domain.	Enables recoverability of NSX Global Manager in the second VMware Cloud Foundation instance if a failure in the first instance occurs.	Must be done manually.
VCF-NSX-GM-RCMD-SEC-001	Establish an operational practice to capture and update the thumbprint of the NSX Local Manager certificate on NSX Global Manager every time the certificate is updated by using SDDC Manager.	<p>Ensures secured connectivity between the NSX Manager instances.</p> <p>Each certificate has its own unique thumbprint. NSX Global Manager stores the unique thumbprint of the NSX Local Manager instances for enhanced security.</p> <p>If an authentication failure between NSX Global Manager and NSX Local Manager occurs, objects</p>	The administrator must establish and follow an operational practice by using a runbook or automated process to ensure that the thumbprint is up-to-date.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		that are created from NSX Global Manager will not be propagated on to the SDN.	

Table 129: NSX Global Manager Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-GM-RCMD-CFG-008	Add the NSX Global Manager appliances to the virtual machine group for the first availability zone.	Ensures that, by default, the NSX Global Manager appliances are powered on a host in the primary availability zone.	Done automatically by VMware Cloud Foundation when stretching a cluster.

NSX Edge Design Elements

Table 130: NSX Edge Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-001	Connect the management interface of each NSX Edge node to the VM management network.	Provides connection from the NSX Manager cluster to the NSX Edge.	None.
VCF-NSX-EDGE-REQD-CFG-002	<ul style="list-style-type: none"> Use two physical NICs as uplinks for the NSX Edge appliances, for example, vmnic0 and vmnic1. Connect the <code>fp-eth0</code> interface of each NSX Edge appliance to a VLAN trunk port group pinned to the one physical NIC (vmnic0) of the host, with the ability to failover to another physical NIC (vmnic1). Connect the <code>fp-eth1</code> interface of each NSX Edge appliance to a VLAN trunk port group pinned to a physical NIC of the host (vmnic1), with the ability to failover to first physical NIC (vmnic0). Leave the additional <code>fp-eth2</code> and <code>fp-eth3</code> 	<ul style="list-style-type: none"> Because VLAN trunk port groups pass traffic for all VLANs, VLAN tagging can occur in the NSX Edge node itself for easy post-deployment configuration. By using two separate VLAN trunk port groups, you can direct traffic from the edge node to a particular host network interface and top of rack switch as needed. In the event of failure of the top of rack switch, the VLAN trunk port group will failover to the other physical NIC and to ensure both <code>fp-eth0</code> and <code>fp-eth1</code> are available. 	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	interfaces of each NSX Edge appliance unused.		
VCF-NSX-EDGE-REQD-CFG-003	Use a dedicated VLAN for edge overlay that is different from the host overlay VLAN.	A dedicated edge overlay network provides support for edge mobility in support of advanced deployments such as multiple availability zones or multi-rack clusters.	<ul style="list-style-type: none"> You must have routing between the VLANs for edge overlay and host overlay. You must allocate another VLAN in the data center infrastructure for edge overlay.
VCF-NSX-EDGE-REQD-CFG-004	<p>Create one uplink profile for the edge nodes with three teaming policies.</p> <ul style="list-style-type: none"> Default teaming policy of load balance source with both active uplinks <code>uplink1</code> and <code>uplink2</code>. Named teaming policy of failover order with a single active uplink <code>uplink1</code> without standby uplinks. Named teaming policy of failover order with a single active uplink <code>uplink2</code> without standby uplinks. 	<ul style="list-style-type: none"> An NSX Edge node that uses a single N-VDS can have only one uplink profile. For increased resiliency and performance, supports the concurrent use of both edge uplinks through both physical NICs on the ESXi hosts. The default teaming policy increases overlay performance and availability by using multiple TEPs, and balancing of overlay traffic. By using named teaming policies, you can connect an edge uplink to a specific host uplink and from there to a specific top of rack switch in the data center. Enables ECMP because the NSX Edge nodes can uplink to the physical network over two different VLANs. 	None.

Table 131: NSX Multi-Rack Edge Availability Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-MR-EDGE-REQD-ENV-001	Place at least one edge VM per edge cluster onto each vSphere cluster, that is in a separate rack.	Provides availability if a failure of a rack or a vSphere cluster occurs.	Additional compute resources required.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-MR-EDGE-REQD-CFG-001	Connect the management interface of each NSX Edge node in each rack to the VM management network that rack.	Provides connection from the NSX Manager cluster to the NSX Edge.	None.
VCF-NSX-MR-EDGE-REQD-CFG-002	Use a dedicated IP pool for the edge overlay network for each rack used for the deployment of a multi-rack NSX Edge cluster.	A dedicated edge overlay network and subnet per rack provides support for a Layer 3 leaf-and-spine physical network design with the Layer 2 boundary at the leaf nodes.	You must manage additional IP pools for each additional rack.
VCF-NSX-MR-EDGE-REQD-CFG-003	Create an edge uplink profile for each rack.	Enables having a dedicated edge overlay VLAN per rack to provide support for a Layer 3 leaf-and-spine physical network design with the Layer 2 boundary at the leaf nodes.	None.

Table 132: NSX Edge Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-EDGE-REQD-CFG-005	Allocate a separate VLAN for edge RTEP overlay that is different from the edge overlay VLAN.	The RTEP network must be on a VLAN that is different from the edge overlay VLAN. This is an NSX requirement that provides support for configuring different MTU size per network.	You must allocate another VLAN in the data center infrastructure.

Table 133: NSX Edge Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-001	Use appropriately sized NSX Edge virtual appliances.	Ensures resource availability and usage efficiency per workload domain.	You must provide sufficient compute resources to support the chosen appliance size.
VCF-NSX-EDGE-RCMD-CFG-002	Deploy the NSX Edge virtual appliances to the default vSphere cluster of the workload domain, sharing the cluster between the workloads and the edge appliances.	Simplifies the configuration and minimizes the number of ESXi hosts required for initial deployment.	Workloads and NSX Edges share the same compute resources.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-003	Deploy two NSX Edge appliances in an edge cluster in the default vSphere cluster of the workload domain.	Creates the minimum size NSX Edge cluster while satisfying the requirements for availability.	For a VI workload domain, additional edge appliances might be required to satisfy increased bandwidth requirements.
VCF-NSX-EDGE-RCMD-CFG-004	Apply VM-VM anti-affinity rules for vSphere DRS to the virtual machines of the NSX Edge cluster.	Keeps the NSX Edge nodes running on different ESXi hosts for high availability.	None.
VCF-NSX-EDGE-RCMD-CFG-005	In vSphere HA, set the restart priority policy for each NSX Edge appliance to high.	<ul style="list-style-type: none"> The NSX Edge nodes are part of the north-south data path for overlay segments. vSphere HA restarts the NSX Edge appliances first to minimise the time an edge VM is offline. Setting the restart priority to high reserves highest for future needs. 	If the restart priority for another VM in the cluster is set to highest, the connectivity delays for edge appliances will be longer.
VCF-NSX-EDGE-RCMD-CFG-006	Create an NSX Edge cluster with the default Bidirectional Forwarding Detection (BFD) configuration between the NSX Edge nodes in the cluster.	<ul style="list-style-type: none"> Satisfies the availability requirements by default. Edge nodes must remain available to create services such as NAT, routing to physical networks, and load balancing. 	None.
VCF-NSX-EDGE-RCMD-CFG-007	Use a single N-VDS in the NSX Edge nodes.	<ul style="list-style-type: none"> Simplifies deployment of the edge nodes. The same N-VDS switch design can be used regardless of edge form factor. Supports multiple TEP interfaces in the edge node. vSphere Distributed Switch is not supported in the edge node. 	None.

Table 134: NSX Edge Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implications
VCF-NSX-EDGE-RCMD-CFG-008	Add the NSX Edge appliances to the virtual	Ensures that, by default, the NSX Edge appliances are	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implications
	machine group for the first availability zone.	powered on upon a host in the primary availability zone.	

BGP Routing Design Elements for VMware Cloud Foundation

Table 135: BGP Routing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-001	To enable ECMP between the Tier-0 gateway and the Layer 3 devices (ToR switches or upstream devices), create two VLANs. The ToR switches or upstream Layer 3 devices have an SVI on one of the two VLANS, and each Edge node in the cluster has an interface on each VLAN.	Supports multiple equal-cost routes on the Tier-0 gateway and provides more resiliency and better bandwidth use in the network.	Additional VLANs are required.
VCF-NSX-BGP-REQD-CFG-002	Assign a named teaming policy to the VLAN segments to the Layer 3 device pair.	Pins the VLAN traffic on each segment to its target edge node interface. From there, the traffic is directed to the host physical NIC that is connected to the target top of rack switch.	None.
VCF-NSX-BGP-REQD-CFG-003	Create a VLAN transport zone for edge uplink traffic.	Enables the configuration of VLAN segments on the N-VDS in the edge nodes.	Additional VLAN transport zones might be required if the edge nodes are not connected to the same top of rack switch pair.
VCF-NSX-BGP-REQD-CFG-004	Deploy a Tier-1 gateway and connect it to the Tier-0 gateway.	Creates a two-tier routing architecture. Abstracts the NSX logical components which interact with the physical data center from the logical components which provide SDN services.	A Tier-1 gateway can only be connected to a single Tier-0 gateway. In cases where multiple Tier-0 gateways are required, you must create multiple Tier-1 gateways.
VCF-NSX-BGP-REQD-CFG-005	Deploy a Tier-1 gateway to the NSX Edge cluster.	Enables stateful services, such as load balancers and NAT, for SDDC management components. Because a Tier-1 gateway always works in active-standby mode, the gateway supports stateful services.	None.

Table 136: BGP Routing Design Requirements for NSX Multi-Rack Edge Availability for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-MRE-BGP-REQD-CFG-001	<p>To enable ECMP between the Tier-0 gateway and the Layer 3 devices, such as leaf switches or upstream devices, create two separate uplink VLANs for the edge nodes in each rack.</p> <p>The leaf switches or Layer 3 upstream devices have an SVI on one of the two VLANs for each rack, and each edge node in the rack has an interface on each VLAN.</p>	Supports multiple equal-cost routes on the Tier-0 gateway and improves resiliency and bandwidth use in the network across multiple racks with a pair of leaf switches in each rack.	Additional VLANs are required.
VCF-NSX-MRE-BGP-REQD-CFG-002	Assign a named teaming policy to the VLAN segments to the Layer 3 device pair for each rack.	Pins the VLAN traffic on each segment to its target edge node interface. From there, the traffic is directed to the host physical NIC that is connected to the target leaf switch in the rack.	None.

Table 137: BGP Routing Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-006	Extend the uplink VLANs to the top of rack switches so that the VLANs are stretched between both availability zones.	Because the NSX Edge nodes will fail over between the availability zones, ensures uplink connectivity to the top of rack switches in both availability zones regardless of the zone the NSX Edge nodes are presently in.	You must configure a stretched Layer 2 network between the availability zones by using physical network infrastructure.
VCF-NSX-BGP-REQD-CFG-007	<p>Provide this SVI configuration on the top of the rack switches.</p> <ul style="list-style-type: none"> In the second availability zone, configure the top of rack switches or upstream Layer 3 devices with an SVI on each of the two uplink VLANs. Make the top of rack switch SVI in both availability zones part of a common stretched Layer 	Enables the communication of the NSX Edge nodes to the top of rack switches in both availability zones over the same uplink VLANs.	You must configure a stretched Layer 2 network between the availability zones by using the physical network infrastructure.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	2 network between the availability zones.		
VCF-NSX-BGP-REQD-CFG-008	<p>Provide this VLAN configuration:</p> <ul style="list-style-type: none"> • Use two VLANs to enable ECMP between the Tier-0 gateway and the Layer 3 devices (top of rack switches or Leaf switches). • The ToR switches or upstream Layer 3 devices have an SVI to one of the two VLANs and each NSX Edge node has an interface to each VLAN. 	Supports multiple equal-cost routes on the Tier-0 gateway, and provides more resiliency and better bandwidth use in the network.	<ul style="list-style-type: none"> • Extra VLANs are required. • Requires stretching uplink VLANs between availability zones
VCF-NSX-BGP-REQD-CFG-009	Create an IP prefix list that permits access to route advertisement by any network instead of using the default IP prefix list.	Used in a route map to prepend a path to one or more autonomous system (AS-path prepend) for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.
VCF-NSX-BGP-REQD-CFG-010	Create a route map-out that contains the custom IP prefix list and an AS-path prepend value set to the Tier-0 local AS added twice.	<ul style="list-style-type: none"> • Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. • Ensures that all ingress traffic passes through the first availability zone. 	<p>You must manually create the route map.</p> <p>The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.</p>
VCF-NSX-BGP-REQD-CFG-011	Create an IP prefix list that permits access to route advertisement by network 0.0.0.0/0 instead of using the default IP prefix list.	Used in a route map to configure local-reference on learned default-route for BGP neighbors in the second availability zone.	You must manually create an IP prefix list that is identical to the default one.
VCF-NSX-BGP-REQD-CFG-012	Apply a route map-in that contains the IP prefix list for the default route 0.0.0.0/0 and assign a lower local-preference, for example, 80, to the learned default route and a lower	<ul style="list-style-type: none"> • Used for configuring neighbor relationships with the Layer 3 devices in the second availability zone. 	<p>You must manually create the route map.</p> <p>The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP</p>

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	local-preference, for example, 90 any routes learned.	<ul style="list-style-type: none"> Ensures that all egress traffic passes through the first availability zone. 	neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.
VCF-NSX-BGP-REQD-CFG-013	Configure the neighbors of the second availability zone to use the route maps as In and Out filters respectively.	Makes the path in and out of the second availability zone less preferred because the AS path is longer and the local preference is lower. As a result, all traffic passes through the first zone.	The two NSX Edge nodes will route north-south traffic through the second availability zone only if the connection to their BGP neighbors in the first availability zone is lost, for example, if a failure of the top of the rack switch pair or in the availability zone occurs.

Table 138: BGP Routing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-BGP-REQD-CFG-014	Extend the Tier-0 gateway to the second VMware Cloud Foundation instance.	<ul style="list-style-type: none"> Supports ECMP north-south routing on all nodes in the NSX Edge cluster. Enables support for cross-instance Tier-1 gateways and cross-instance network segments. 	The Tier-0 gateway deployed in the second instance is removed.
VCF-NSX-BGP-REQD-CFG-015	Set the Tier-0 gateway as primary for all VMware Cloud Foundation instances.	<ul style="list-style-type: none"> In NSX Federation, a Tier-0 gateway lets egress traffic from connected Tier-1 gateways only in its primary locations. Local ingress and egress traffic is controlled independently at the Tier-1 level. No segments are provisioned directly to the Tier-0 gateway. A mixture of network spans (local to a VMware Cloud Foundation instance or spanning multiple instances) is enabled without requiring additional Tier-0 	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		<p>gateways and hence edge nodes.</p> <ul style="list-style-type: none"> If a failure in a VMware Cloud Foundation instance occurs, the local-instance networking in the other instance remains available without manual intervention. 	
VCF-NSX-BGP-REQD-CFG-016	From the global Tier-0 gateway, establish BGP neighbor peering to the ToR switches connected to the second VMware Cloud Foundation instance.	<ul style="list-style-type: none"> Enables the learning and advertising of routes in the second VMware Cloud Foundation instance. Facilitates a potential automated failover of networks from the first to the second VMware Cloud Foundation instance. 	None.
VCF-NSX-BGP-REQD-CFG-017	Use a stretched Tier-1 gateway and connect it to the Tier-0 gateway for cross-instance networking.	<ul style="list-style-type: none"> Enables network span between the VMware Cloud Foundation instances because NSX network segments follow the span of the gateway they are attached to. Creates a two-tier routing architecture. 	None.
VCF-NSX-BGP-REQD-CFG-018	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the stretched Tier-1 gateway. Set the first VMware Cloud Foundation instance as primary and the second instance as secondary.	<ul style="list-style-type: none"> Enables cross-instance network span between the first and second VMware Cloud Foundation instances. Enables deterministic ingress and egress traffic for the cross-instance network. If a VMware Cloud Foundation instance failure occurs, enables deterministic failover of the Tier-1 traffic flow. During the recovery of the inaccessible VMware Cloud Foundation instance, enables deterministic failback of the Tier-1 traffic flow, 	You must manually fail over and fail back the cross-instance network from the standby NSX Global Manager.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		<p>preventing unintended asymmetrical routing.</p> <ul style="list-style-type: none"> Eliminates the need to use BGP attributes in the first and second VMware Cloud Foundation instances to influence location preference and failover. 	
VCF-NSX-BGP-REQD-CFG-019	Assign the NSX Edge cluster in each VMware Cloud Foundation instance to the local Tier-1 gateway for that VMware Cloud Foundation instance.	<ul style="list-style-type: none"> Enables instance-specific networks to be isolated to their specific instances. Enables deterministic flow of ingress and egress traffic for the instance-specific networks. 	You can use the service router that is created for the Tier-1 gateway for networking services. However, such configuration is not required for network connectivity.
VCF-NSX-BGP-REQD-CFG-020	Set each local Tier-1 gateway only as primary in that instance. Avoid setting the gateway as secondary in the other instances.	Prevents the need to use BGP attributes in primary and secondary instances to influence the instance ingress-egress preference.	None.

Table 139: BGP Routing Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
VCF-NSX-BGP-RCMD-CFG-001	Deploy an active-active Tier-0 gateway.	Supports ECMP north-south routing on all Edge nodes in the NSX Edge cluster.	Active-active Tier-0 gateways cannot provide stateful services such as NAT.
VCF-NSX-BGP-RCMD-CFG-002	Configure the BGP Keep Alive Timer to 4 and Hold Down Timer to 12 or lower between the top of rack switches and the Tier-0 gateway.	Provides a balance between failure detection between the top of rack switches and the Tier-0 gateway, and overburdening the top of rack switches with keep-alive traffic.	<p>By using longer timers to detect if a router is not responding, the data about such a router remains in the routing table longer. As a result, the active router continues to send traffic to a router that is down.</p> <p>These timers must be aligned with the data center fabric design of your organization.</p>
VCF-NSX-BGP-RCMD-CFG-003	Do not enable Graceful Restart between BGP neighbors.	<p>Avoids loss of traffic.</p> <p>On the Tier-0 gateway, BGP peers from all the gateways are always active. On a failover, the Graceful Restart</p>	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Recommendation Justification	Recommendation Implication
		capability increases the time a remote neighbor takes to select an alternate Tier-0 gateway. As a result, BFD-based convergence is delayed.	
VCF-NSX-BGP-RCMD-CFG-004	Enable helper mode for Graceful Restart mode between BGP neighbors.	Avoids loss of traffic. During a router restart, helper mode works with the graceful restart capability of upstream routers to maintain the forwarding table which in turn will forward packets to a down neighbor even after the BGP timers have expired causing loss of traffic.	None.
VCF-NSX-BGP-RCMD-CFG-005	Enable Inter-SR iBGP routing.	In the event that an edge node has all of its northbound eBGP sessions down, north-south traffic will continue to flow by routing traffic to a different edge node.	None.
VCF-NSX-BGP-RCMD-CFG-006	Deploy a Tier-1 gateway in non-preemptive failover mode.	Ensures that after a failed NSX Edge transport node is back online, it does not take over the gateway services thus preventing a short service outage.	None.
VCF-NSX-BGP-RCMD-CFG-007	Enable standby relocation of the Tier-1 gateway.	Ensures that if an edge failure occurs, a standby Tier-1 gateway is created on another edge node.	None.

Table 140: BGP Routing Design Recommendations for NSX Federation in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-BGP-RCMD-CFG-008	Use Tier-1 gateways to control the span of networks and ingress and egress traffic in the VMware Cloud Foundation instances.	Enables a mixture of network spans (isolated to a VMware Cloud Foundation instance or spanning multiple instances) without requiring additional Tier-0 gateways and hence edge nodes.	To control location span, a Tier-1 gateway must be assigned to an edge cluster and hence has the Tier-1 SR component. East-west traffic between Tier-1 gateways with SRs need to physically traverse an edge node.
VCF-NSX-BGP-RCMD-CFG-009	Allocate a Tier-1 gateway in each instance for instance-specific networks and	<ul style="list-style-type: none"> Creates a two-tier routing architecture. 	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
	connect it to the stretched Tier-0 gateway.	<ul style="list-style-type: none"> Enables local-instance networks that are not to span between the VMware Cloud Foundation instances. Guarantees that local-instance networks remain available if a failure occurs in another VMware Cloud Foundation instance. 	

Overlay Design Elements for VMware Cloud Foundation

Table 141: Overlay Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-OVERLAY-REQD-CFG-001	Configure all ESXi hosts in the workload domain as transport nodes in NSX.	Enables distributed routing, logical segments, and distributed firewall.	None.
VCF-NSX-OVERLAY-REQD-CFG-002	Configure each ESXi host as a transport node using transport node profiles.	<ul style="list-style-type: none"> Enables the participation of ESXi hosts and the virtual machines running on them in NSX overlay and VLAN networks. Transport node profiles can only be applied at the cluster level. 	None.
VCF-NSX-OVERLAY-REQD-CFG-003	To provide virtualized network capabilities to workloads, use overlay networks with NSX Edge nodes and distributed routing.	<ul style="list-style-type: none"> Creates isolated, multi-tenant broadcast domains across data center fabrics to deploy elastic, logical networks that span physical network boundaries. Enables advanced deployment topologies by introducing Layer 2 abstraction from the data center networks. 	Requires configuring transport networks with an MTU size of at least 1,600 bytes.
VCF-NSX-OVERLAY-REQD-CFG-004	Create a single overlay transport zone in the NSX instance for all overlay traffic across the host and NSX Edge transport nodes of the workload domain or multiple workload domains using a shared NSX instance.	<ul style="list-style-type: none"> Ensures that overlay segments are connected to an NSX Edge node for services and north-south routing. Ensures that all segments are available to all ESXi 	All clusters in all workload domains that share the same NSX Manager share the same overlay transport zone.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		hosts and NSX Edge nodes configured as transport nodes.	
VCF-NSX-OVERLAY-REQD-CFG-005	Create an uplink profile with a load balance source teaming policy with two active uplinks for ESXi hosts.	For increased resiliency and performance, supports the concurrent use of both physical NICs on the ESXi hosts that are configured as transport nodes.	None.

Table 142: Overlay Design Requirements for a Multi-Rack Compute VI Workload Domain Cluster for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-L3MR-OVERLAY-REQD-CFG-001	Create an uplink profile for each rack with a separate host TEP transport VLAN ID per rack.	Enables a Layer 3 boundary at the top-of-rack switches for host TEP traffic.	Requires additional host TEP VLAN ID per rack.
VCF-NSX-L3MR-OVERLAY-REQD-CFG-002	Create a host TEP IP pool for each rack with a subnet allocated per rack and a gateway for the subnet.	<ul style="list-style-type: none"> Provides the host TEP IP addressing for each rack Enables a Layer 3 boundary at the top-of-rack switches for host TEP traffic. Provides routing capabilities for the host TEPs between racks. 	Requires additional subnet per rack.
VCF-NSX-L3MR-OVERLAY-REQD-CFG-003	Create NSX host sub transport node profiles for additional racks.	<ul style="list-style-type: none"> Enables NSX host transport node configuration per rack. 	None

Table 143: Overlay Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-OVERLAY-RCMD-CFG-001	Use static IP pools to assign IP addresses to the host TEP interfaces.	<ul style="list-style-type: none"> Removes the need for an external DHCP server for the host overlay VLANs. You can use NSX Manager to verify static IP pool configurations. 	None.
VCF-NSX-OVERLAY-RCMD-CFG-002	Use hierarchical two-tier replication on all overlay segments.	Hierarchical two-tier replication is more efficient because it reduced the number of ESXi hosts the source ESXi host must replicate traffic to if hosts	None.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		have different TEP subnets. This is typically the case with more than one cluster and will improve performance in that scenario.	

Table 144: Overlay Design Recommendations for Stretched Clusters in VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-OVERLAY-RCMD-CFG-003	Configure an NSX sub-transport node profile.	<ul style="list-style-type: none"> You can use static IP pools for the host TEPs in each availability zone. The NSX transport node profile can remain attached when using two separate VLANs for host TEPs at each availability zone as required for clusters that are based on vSphere Lifecycle Manager images. Using an external DHCP server for the host overlay VLANs in both availability zones is not required. 	Changes to the host transport node configuration are done at the vSphere cluster level.

Application Virtual Network Design Elements for VMware Cloud Foundation

Table 145: Application Virtual Network Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-001	Create one cross-instance NSX segment for the components of a VMware Aria Suite application or another solution that requires mobility between VMware Cloud Foundation instances.	<p>Prepares the environment for the deployment of solutions on top of VMware Cloud Foundation, such as VMware Aria Suite, without a complex physical network configuration.</p> <p>The components of the VMware Aria Suite application must be easily portable between VMware Cloud Foundation instances without requiring reconfiguration.</p>	Each NSX segment requires a unique IP address space.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-002	Create one or more local-instance NSX segments for the components of a VMware Aria Suite application or another solution that are assigned to a specific VMware Cloud Foundation instance.	Prepares the environment for the deployment of solutions on top of VMware Cloud Foundation, such as VMware Aria Suite, without a complex physical network configuration.	Each NSX segment requires a unique IP address space.

Table 146: Application Virtual Network Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-AVN-REQD-CFG-003	Extend the cross-instance NSX segment to the second VMware Cloud Foundation instance.	Enables workload mobility without a complex physical network configuration. The components of a VMware Aria Suite application must be easily portable between VMware Cloud Foundation instances without requiring reconfiguration.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-004	In each VMware Cloud Foundation instance, create additional local-instance NSX segments.	Enables workload mobility within a VMware Cloud Foundation instance without complex physical network configuration. Each VMware Cloud Foundation instance should have network segments to support workloads which are isolated to that VMware Cloud Foundation instance.	Each NSX segment requires a unique IP address space.
VCF-NSX-AVN-REQD-CFG-005	In each VMware Cloud Foundation instance, connect or migrate the local-instance NSX segments to the corresponding local-instance Tier-1 gateway.	Configures local-instance NSX segments at required sites only.	Requires an individual Tier-1 gateway for local-instance segments.

Table 147: Application Virtual Network Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-NSX-AVN-RCMD-CFG-001	Use overlay-backed NSX segments.	<ul style="list-style-type: none"> Supports expansion to deployment topologies for 	Using overlay-backed NSX segments requires routing, eBGP recommended,

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		multiple VMware Cloud Foundation instances. <ul style="list-style-type: none"> Limits the number of VLANs required for the data center fabric. 	between the data center fabric and edge nodes.

Load Balancing Design Elements for VMware Cloud Foundation

Table 148: Load Balancing Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-001	Deploy a standalone Tier-1 gateway to support advanced stateful services such as load balancing for other management components.	Provides independence between north-south Tier-1 gateways to support advanced deployment scenarios.	You must add a separate Tier-1 gateway.
VCF-NSX-LB-REQD-CFG-002	When creating load balancing services for Application Virtual Networks, connect the standalone Tier-1 gateway to the cross-instance NSX segments.	Provides load balancing to applications connected to the cross-instance network.	You must connect the gateway to each network that requires load balancing.
VCF-NSX-LB-REQD-CFG-003	Configure a default static route on the standalone Tier-1 gateway with a next hop the Tier-1 gateway for the segment to provide connectivity to the load balancer.	Because the Tier-1 gateway is standalone, it does not auto-configure its routes.	None.

Table 149: Load Balancing Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-NSX-LB-REQD-CFG-004	Deploy a standalone Tier-1 gateway in the second VMware Cloud Foundation instance.	Provides a cold-standby non-global service router instance for the second VMware Cloud Foundation instance to support services on the cross-instance network which require advanced services not currently supported as NSX global objects.	<ul style="list-style-type: none"> You must add a separate Tier-1 gateway. You must manually configure any services and synchronize them between the non-global service router instances in the first and second VMware Cloud Foundation instances. To avoid a network conflict between the two VMware Cloud

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
			Foundation instances, make sure that the primary and standby networking services are not both active at the same time.
VCF-NSX-LB-REQD-CFG-005	Connect the standalone Tier-1 gateway in the second VMware Cloud Foundation instance to the cross-instance NSX segment.	Provides load balancing to applications connected to the cross-instance network in the second VMware Cloud Foundation instance.	You must connect the gateway to each network that requires load balancing.
VCF-NSX-LB-REQD-CFG-006	Configure a default static route on the standalone Tier-1 gateway in the second VMware Cloud Foundation instance with a next hop as the Tier-1 gateway for the segment it connects with to provide connectivity to the load balancers.	Because the Tier-1 gateway is standalone, it does not autoconfigure its routes.	None.
VCF-NSX-LB-REQD-CFG-007	Establish a process to ensure any changes made on to the load balancer instance in the first VMware Cloud Foundation instance are manually applied to the disconnected load balancer in the second instance.	<p>Keeps the network service in the failover load balancer instance ready for activation if a failure in the first VMware Cloud Foundation instance occurs.</p> <p>Because network services are not supported as global objects, you must configure them manually in each VMware Cloud Foundation instance. The load balancer service in one instance must be connected and active, while the service in the other instance must be disconnected and inactive.</p>	<ul style="list-style-type: none"> • Because of incorrect configuration between the VMware Cloud Foundation instances, the load balancer service in the second instance might come online with an invalid or incomplete configuration. • If both VMware Cloud Foundation instances are online and active at the same time, a conflict between services could occur resulting in a potential outage. • The administrator must establish and follow an operational practice by using a runbook or automated process to ensure that configuration changes are reproduced in each VMware Cloud Foundation instance.

SDDC Manager Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to SDDC Manager in an environment with a single or multiple VMware Cloud Foundation instances.

For full design details, see [SDDC Manager Design for](#) .

Table 150: SDDC Manager Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-SDDCMGR-REQD-CFG-001	Deploy an SDDC Manager system in the first availability zone of the management domain.	SDDC Manager is required to perform VMware Cloud Foundation capabilities, such as provisioning VI workload domains, deploying solutions, patching, upgrading, and others.	None.
VCF-SDDCMGR-REQD-CFG-002	Deploy SDDC Manager with its default configuration.	The configuration of SDDC Manager is not configurable and should not be changed from its defaults.	None.
VCF-SDDCMGR-REQD-CFG-003	Place the SDDC Manager appliance on the VM management network.	<ul style="list-style-type: none"> Simplifies IP addressing for management VMs by using the same VLAN and subnet. Provides simplified secure access to management VMs in the same VLAN network. 	None.

Table 151: SDDC Manager Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDCMGR-RCMD-CFG-001	Connect SDDC Manager to the Internet for downloading software bundles.	SDDC Manager must be able to download install and upgrade software bundles for deployment of VI workload domains and solutions, and for upgrade from a repository.	The rules of your organization might not permit direct access to the Internet. In this case, you must either download software bundles for SDDC Manager manually, or configure an offline depot.
VCF-SDDCMGR-RCMD-CFG-002	Configure an authenticated network proxy to connect SDDC Manager to the Internet.	To protect SDDC Manager against external attacks from the Internet.	You must manage the proxy settings and security manually.
VCF-SDDCMGR-RCMD-CFG-003	Configure SDDC Manager with a VMware Customer Connect account with VMware Cloud Foundation	Software bundles for VMware Cloud Foundation are stored in a repository that is secured with access controls.	Requires the use of a VMware Customer Connect user account with access to

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
	entitlement to check for and download software bundles.		VMware Cloud Foundation licensing. Sites without an internet connection can use local upload option instead.
VCF-SDDCMGR-RCMD-CFG-004	Configure SDDC Manager with an external certificate authority that is responsible for providing signed certificates.	Provides increased security by implementing signed certificate generation and replacement across the management components.	An external certificate authority, such as Microsoft CA, must be locally available.

VMware Aria Suite Lifecycle Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to VMware Aria Suite Lifecycle in an environment with a single or multiple VMware Cloud Foundation instances.

For full design details, see [Design for](#) .

Table 152: VMware Aria Suite Lifecycle Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-001	Deploy a VMware Aria Suite Lifecycle instance in the management domain of each VMware Cloud Foundation instance to provide life cycle management for VMware Aria Suite and Workspace ONE Access.	Provides life cycle management operations for VMware Aria Suite applications and Workspace ONE Access.	You must ensure that the required resources are available.
VCF-VASL-REQD-CFG-002	Deploy VMware Aria Suite Lifecycle by using SDDC Manager.	<ul style="list-style-type: none"> Deploys VMware Aria Suite Lifecycle in VMware Cloud Foundation mode, which enables the integration with the SDDC Manager inventory for product deployment and life cycle management of VMware Aria Suite components. Automatically configures the standalone Tier-1 gateway required for load balancing the clustered Workspace ONE Access and VMware Aria Suite components. 	None.
VCF-VASL-REQD-CFG-003	Allocate extra 100 GB of storage to the VMware Aria Suite Lifecycle appliance for VMware Aria Suite product binaries.	<ul style="list-style-type: none"> Provides support for VMware Aria Suite product binaries (install, upgrade, and patch) and content management. 	None.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
		<ul style="list-style-type: none"> SDDC Manager automates the creation of storage. 	
VCF-VASL-REQD-CFG-004	Place the VMware Aria Suite Lifecycle appliance on an overlay-backed (recommended) or VLAN-backed NSX network segment.	Provides a consistent deployment model for management applications.	You must use an implementation in NSX to support this networking configuration.
VCF-VASL-REQD-CFG-005	Import VMware Aria Suite product licenses to the Locker repository for product life cycle operations.	<ul style="list-style-type: none"> You can review the validity, details, and deployment usage for the license across the VMware Aria Suite products. You can reference and use licenses during product life cycle operations, such as deployment and license replacement. 	When using the API, you must specify the Locker ID for the license to be used in the JSON payload.
VCF-VASL-REQD-ENV-001	Configure datacenter objects in VMware Aria Suite Lifecycle for local and cross-instance VMware Aria Suite deployments and assign the management domain vCenter Server instance to each data center.	You can deploy and manage the integrated VMware Aria Suite components across the SDDC as a group.	You must manage a separate datacenter object for the products that are specific to each instance.
VCF-VASL-REQD-ENV-002	If deploying VMware Aria Operations for Logs, create a local-instance environment in VMware Aria Suite Lifecycle.	Supports the deployment of an instance of VMware Aria Operations for Logs.	None.
VCF-VASL-REQD-ENV-003	If deploying VMware Aria Operations or VMware Aria Automation, create a cross-instance environment in VMware Aria Suite Lifecycle	<ul style="list-style-type: none"> Supports deployment and management of the integrated VMware Aria Suite products across VMware Cloud Foundation instances as a group. Enables the deployment of instance-specific components, such as VMware Aria Operations remote collectors. In VMware Aria Suite Lifecycle, you can deploy and manage VMware Aria Operations remote collector objects only in an environment that contains the associated cross-instance components. 	You can manage instance-specific components, such as remote collectors, only in an environment that is cross-instance.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-SEC-001	Use the custom vCenter Server role for VMware Aria Suite Lifecycle that has the minimum privileges required to support the deployment and upgrade of VMware Aria Suite products.	VMware Aria Suite Lifecycle accesses vSphere with the minimum set of permissions that are required to support the deployment and upgrade of VMware Aria Suite products. SDDC Manager automates the creation of the custom role.	You must maintain the permissions required by the custom role.
VCF-VASL-REQD-SEC-002	Use the service account in vCenter Server for application-to-application communication from VMware Aria Suite Lifecycle to vSphere. Assign global permissions using the custom role.	<ul style="list-style-type: none"> Provides the following access control features: <ul style="list-style-type: none"> VMware Aria Suite Lifecycle accesses vSphere with the minimum set of required permissions. You can introduce improved accountability in tracking request-response interactions between the components of the SDDC. SDDC Manager automates the creation of the service account. 	<ul style="list-style-type: none"> You must maintain the life cycle and availability of the service account outside of SDDC manager password rotation.

Table 153: VMware Aria Suite Lifecycle Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-006	For multiple availability zones, add the VMware Aria Suite Lifecycle appliance to the VM group for the first availability zone.	Ensures that, by default, the VMware Aria Suite Lifecycle appliance is powered on a host in the first availability zone.	If VMware Aria Suite Lifecycle is deployed after the creation of the stretched management cluster, you must add the VMware Aria Suite Lifecycle appliance to the VM group manually.

Table 154: VMware Aria Suite Lifecycle Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-007	Configure the DNS settings for the VMware Aria Suite Lifecycle appliance to use DNS servers in each instance.	Improves resiliency in the event of an outage of external services for a VMware Cloud Foundation instance.	As you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the DNS settings of the VMware Aria Suite Lifecycle appliance must be updated.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
VCF-VASL-REQD-CFG-008	Configure the NTP settings for the VMware Aria Suite Lifecycle appliance to use NTP servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	As you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the NTP settings on the VMware Aria Suite Lifecycle appliance must be updated.
VCF-VASL-REQD-ENV-004	Assign the management domain vCenter Server instance in the additional VMware Cloud Foundation instance to the cross-instance data center.	Supports the deployment of VMware Aria Operations remote collectors in an additional VMware Cloud Foundation instance.	None.

Table 155: VMware Aria Suite Lifecycle Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-VASL-RCMD-CFG-001	Protect VMware Aria Suite Lifecycle by using vSphere HA.	Supports the availability objectives for VMware Aria Suite Lifecycle without requiring manual intervention during a failure event.	None.
VCF-VASL-RCMD-LCM-001	Obtain product binaries for install, patch, and upgrade in VMware Aria Suite Lifecycle from VMware Customer Connect.	<ul style="list-style-type: none"> You can upgrade VMware Aria Suite products based on their general availability and endpoint interoperability rather than being listed as part of VMware Cloud Foundation bill of materials (BOM). You can deploy and manage binaries in an environment that does not allow access to the Internet or are dark sites. 	<p>The site must have an Internet connection to use VMware Customer Connect.</p> <p>Sites without an Internet connection should use the local upload option instead.</p>
VCF-VASL-RCMD-LCM-002	Use support packs (PSPAKS) for VMware Aria Suite Lifecycle to enable upgrading to later versions of VMware Aria Suite products.	Enables the upgrade of an existing VMware Aria Suite Lifecycle to permit later versions of VMware Aria Suite products without an associated VMware Cloud Foundation upgrade. See VMware Knowledge Base article 88829	None.
VCF-VASL-RCMD-SEC-001	Enable integration between VMware Aria Suite Lifecycle and your corporate identity source by using the Workspace ONE Access instance.	<ul style="list-style-type: none"> Enables authentication to VMware Aria Suite Lifecycle by using your corporate identity source. Enables authorization through the assignment of 	You must deploy and configure Workspace ONE Access to establish the integration between VMware Aria Suite Lifecycle and your corporate identity sources.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		organization and cloud services roles to enterprise users and groups defined in your corporate identity source.	
VCF-VASL-RCMD-SEC-002	Create corresponding security groups in your corporate directory services for VMware Aria Suite Lifecycle roles: <ul style="list-style-type: none"> • VCF • Content Release Manager • Content Developer 	Streamlines the management of VMware Aria Suite Lifecycle roles for users.	<ul style="list-style-type: none"> • You must create the security groups outside of the SDDC stack. • You must set the desired directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period.

Workspace ONE Access Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to Workspace ONE Access in an environment with a single or multiple VMware Cloud Foundation instances. The design elements also considers whether the management domain has a single or multiple availability zones.

For full design details, see [Design for](#) .

Table 156: Workspace ONE Access Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-ENV-001	Create a global environment in VMware Aria Suite Lifecycle to support the deployment of Workspace ONE Access.	A global environment is required by VMware Aria Suite Lifecycle to deploy Workspace ONE Access.	None.
VCF-WSA-REQD-SEC-001	Import certificate authority-signed certificates to the Locker repository for Workspace ONE Access product life cycle operations.	<ul style="list-style-type: none"> • You can reference and use certificate authority-signed certificates during product life cycle operations, such as deployment and certificate replacement. 	When using the API, you must specify the Locker ID for the certificate to be used in the JSON payload.
VCF-WSA-REQD-CFG-001	Deploy an appropriately sized Workspace ONE Access instance according to the deployment model you have selected by using VMware Aria Suite Lifecycle in VMware Cloud Foundation mode.	The Workspace ONE Access instance is managed by VMware Aria Suite Lifecycle and imported into the SDDC Manager inventory.	None.
VCF-WSA-REQD-CFG-002	Place the Workspace ONE Access appliances on an	Provides a consistent deployment model for	You must use an implementation in NSX to

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	overlay-backed or VLAN-backed NSX network segment.	management applications in an environment with a single or multiple VMware Cloud Foundation instances.	support this network configuration.
VCF-WSA-REQD-CFG-003	Use the embedded PostgreSQL database with Workspace ONE Access.	Removes the need for external database services.	None.
VCF-WSA-REQD-CFG-004	Add a VM group for Workspace ONE Access and set VM rules to restart the Workspace ONE Access VM group before any of the VMs that depend on it for authentication.	You can define the startup order of virtual machines regarding the service dependency. The startup order ensures that vSphere HA powers on the Workspace ONE Access virtual machines in an order that respects product dependencies.	None.
VCF-WSA-REQD-CFG-005	Connect the Workspace ONE Access instance to a supported upstream Identity Provider.	You can integrate your enterprise directory with Workspace ONE Access to synchronize users and groups to the Workspace ONE Access identity and access management services.	None.
VCF-WSA-REQD-CFG-006	If using clustered Workspace ONE Access, configure second and third native connectors that correspond to the second and third Workspace ONE Access cluster nodes to support the high availability of directory services access.	Adding the additional native connectors provides redundancy and improves performance by load-balancing authentication requests.	Each of the Workspace ONE Access cluster nodes must be joined to the Active Directory domain to use Active Directory with Integrated Windows Authentication with the native connector.
VCF-WSA-REQD-CFG-007	If using clustered Workspace ONE Access, use the NSX load balancer that is configured by SDDC Manager on a dedicated Tier-1 gateway.	<ul style="list-style-type: none"> During the deployment of Workspace ONE Access by using VMware Aria Suite Lifecycle, SDDC Manager automates the configuration of an NSX load balancer for Workspace ONE Access to facilitate scale-out. 	You must use the load balancer that is configured by SDDC Manager and the integration with VMware Aria Suite Lifecycle.

Table 157: Workspace ONE Access Design Requirements for Stretched Clusters in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-008	Add the Workspace ONE Access appliances to the	Ensures that, by default, the Workspace ONE Access	<ul style="list-style-type: none"> If the Workspace ONE Access instance is

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	VM group for the first availability zone.	cluster nodes are powered on a host in the first availability zone.	<p>deployed after the creation of the stretched management cluster, you must add the appliances to the VM group manually.</p> <ul style="list-style-type: none"> Clustered Workspace ONE Access might require manual intervention after a failure of the active availability zone occurs.

Table 158: Workspace ONE Access Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-WSA-REQD-CFG-009	Configure the DNS settings for Workspace ONE Access to use DNS servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	None.
VCF-WSA-REQD-CFG-010	Configure the NTP settings on Workspace ONE Access cluster nodes to use NTP servers in each VMware Cloud Foundation instance.	Improves resiliency if an outage of external services for a VMware Cloud Foundation instance occurs.	If you scale from a deployment with a single VMware Cloud Foundation instance to one with multiple VMware Cloud Foundation instances, the NTP settings on Workspace ONE Access must be updated.

Table 159: Workspace ONE Access Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-WSA-RCMD-CFG-001	Protect all Workspace ONE Access nodes using vSphere HA.	Supports high availability for Workspace ONE Access.	<p>None for standard deployments.</p> <p>Clustered Workspace ONE Access deployments might require intervention if an ESXi host failure occurs.</p>
VCF-WSA-RCMD-CFG-002	When using Active Directory as an Identity Provider, use Active Directory over LDAP as the Directory Service connection option.	The native (embedded) Workspace ONE Access connector binds to Active Directory over LDAP using a standard bind authentication.	<ul style="list-style-type: none"> In a multi-domain forest, where the Workspace ONE Access instance connects to a child domain, Active Directory security groups must have global scope. Therefore, members

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
			<p>added to the Active Directory global security group must reside within the same Active Directory domain.</p> <ul style="list-style-type: none"> If authentication to more than one Active Directory domain is required, additional Workspace ONE Access directories are required.
VCF-WSA-RCMD-CFG-003	When using Active Directory as an Identity Provider, use an Active Directory user account with a minimum of read-only access to Base DNs for users and groups as the service account for the Active Directory bind.	<p>Provides the following access control features:</p> <ul style="list-style-type: none"> Workspace ONE Access connects to the Active Directory with the minimum set of required permissions to bind and query the directory. You can introduce improved accountability in tracking request-response interactions between the Workspace ONE Access and Active Directory. 	<ul style="list-style-type: none"> You must manage the password life cycle of this account. If authentication to more than one Active Directory domain is required, additional accounts are required for the Workspace ONE Access connector to bind to each Active Directory domain over LDAP.
VCF-WSA-RCMD-CFG-004	Configure the directory synchronization to synchronize only groups required for the integrated SDDC solutions.	<ul style="list-style-type: none"> Limits the number of replicated groups required for each product. Reduces the replication interval for group information. 	You must manage the groups from your enterprise directory selected for synchronization to Workspace ONE Access.
VCF-WSA-RCMD-CFG-005	Activate the synchronization of enterprise directory group members when a group is added to the Workspace ONE Access directory.	When activated, members of the enterprise directory groups are synchronized to the Workspace ONE Access directory when groups are added. When deactivated, group names are synchronized to the directory, but members of the group are not synchronized until the group is entitled to an application or the group name is added to an access policy.	None.
VCF-WSA-RCMD-CFG-006	Enable Workspace ONE Access to synchronize	Allows Workspace ONE Access to update and cache the membership of groups	Changes to group membership are not

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
	nested group members by default.	without querying your enterprise directory.	reflected until the next synchronization event.
VCF-WSA-RCMD-CFG-007	Add a filter to the Workspace ONE Access directory settings to exclude users from the directory replication.	Limits the number of replicated users for Workspace ONE Access within the maximum scale.	To ensure that replicated user accounts are managed within the maximums, you must define a filtering schema that works for your organization based on your directory attributes.
VCF-WSA-RCMD-CFG-008	Configure the mapped attributes included when a user is added to the Workspace ONE Access directory.	You can configure the minimum required and extended user attributes to synchronize directory user accounts for the Workspace ONE Access to be used as an authentication source for cross-instance VMware Aria Suite solutions.	<p>User accounts in your organization's enterprise directory must have the following required attributes mapped:</p> <ul style="list-style-type: none"> • <code>firstname</code>, for example, <code>givenname</code> for Active Directory • <code>lastName</code>, for example, <code>sn</code> for Active Directory • <code>email</code>, for example, <code>mail</code> for Active Directory • <code>userName</code>, for example, <code>sAMAccountName</code> for Active Directory • If you require users to sign in with an alternate unique identifier, for example, <code>userPrincipalName</code>, you must map the attribute and update the identity and access management preferences.
VCF-WSA-RCMD-CFG-009	Configure the Workspace ONE Access directory synchronization frequency to a reoccurring schedule, for example, 15 minutes.	Ensures that any changes to group memberships in the corporate directory are available for integrated solutions in a timely manner.	Schedule the synchronization interval to be longer than the time to synchronize from the enterprise directory. If users and groups are being synchronized to Workspace ONE Access when the next synchronization is scheduled, the new synchronization starts immediately after the end of the previous iteration. With

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
			this schedule, the process is continuous.
VCF-WSA-RCMD-SEC-001	Create corresponding security groups in your corporate directory services for these Workspace ONE Access roles: <ul style="list-style-type: none"> • Super Admin • Directory Admins • ReadOnly Admin 	Streamlines the management of Workspace ONE Access roles to users.	<ul style="list-style-type: none"> • You must set the appropriate directory synchronization interval in Workspace ONE Access to ensure that changes are available within a reasonable period. • You must create the security group outside of the SDDC stack.
VCF-WSA-RCMD-SEC-002	Configure a password policy for Workspace ONE Access local directory users, admin and configadmin .	<p>You can set a policy for Workspace ONE Access local directory users that addresses your corporate policies and regulatory standards.</p> <p>The password policy is applicable only to the local directory users and does not impact your organization directory.</p>	<p>You must set the policy in accordance with your organization policies and regulatory standards, as applicable.</p> <p>You must apply the password policy on the Workspace ONE Access cluster nodes.</p>

Lifecycle Management Design Elements for VMware Cloud Foundation

Use this list of requirements for reference related to life cycle management in a VMware Cloud Foundation environment.

For full design details, see [Lifecycle Management Design for VMware Cloud Foundation](#).

Table 160: Lifecycle Management Design Requirements for VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-001	Use SDDC Manager to perform the life cycle management of the following components: <ul style="list-style-type: none"> • SDDC Manager • NSX Manager • NSX Edges • vCenter Server • ESXi 	Because the deployment scope of SDDC Manager covers the full VMware Cloud Foundation stack, SDDC Manager performs patching, update, or upgrade of these components across all workload domains.	The operations team must understand and be aware of the impact of a patch, update, or upgrade operation by using SDDC Manager.
VCF-LCM-REQD-002	Use VMware Aria Suite Lifecycle to manage the life	VMware Aria Suite Lifecycle automates the life cycle of VMware Aria Suite Lifecycle	<ul style="list-style-type: none"> • You must deploy VMware Aria Suite Lifecycle by using SDDC Manager.

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
	<p>cycle of the following components:</p> <ul style="list-style-type: none"> VMware Aria Suite Lifecycle Workspace ONE Access 	and Workspace ONE Access.	<ul style="list-style-type: none"> You must manually apply Workspace ONE Access patches, updates, and hotfixes. Patches, updates, and hotfixes for Workspace ONE Access are not generally managed by VMware Aria Suite Lifecycle.

Table 161: Lifecycle Management Design Recommendations for VMware Cloud Foundation

Requirement ID	Design Recommendations	Justification	Implication
VCF-LCM-RCMD-001	Use vSphere Lifecycle Manager images to manage the life cycle of vSphere clusters.	<ul style="list-style-type: none"> With vSphere Lifecycle Manager images, firmware updates are carried out through firmware and driver add-ons, which you add to the image you use to manage a cluster. You can check the hardware compatibility of the hosts in a cluster against the VMware Compatibility Guide. You can validate a vSphere Lifecycle Manager image to check if it applies to all hosts in the cluster. You can also perform a remediation pre-check. 	<ul style="list-style-type: none"> Updating the firmware with images requires an OEM-provided hardware support manager plug-in, which integrates with vSphere Lifecycle Manager. An updated vSAN Hardware Compatibility List (vSAN HCL) is required during bring-up.

Table 162: Lifecycle Management Design Requirements for NSX Federation in VMware Cloud Foundation

Requirement ID	Design Requirement	Justification	Implication
VCF-LCM-REQD-003	Use the upgrade coordinator in NSX to perform life cycle management on the NSX Global Manager appliances.	The version of SDDC Manager in this design is not currently capable of life cycle operations (patching, update, or upgrade) for NSX Global Manager.	<ul style="list-style-type: none"> You must explicitly plan upgrades of the NSX Global Manager nodes. An upgrade of the NSX Global Manager nodes might require a cascading upgrade of the NSX Local Manager nodes and underlying SDDC Manager infrastructure

Table continued on next page

Continued from previous page

Requirement ID	Design Requirement	Justification	Implication
			before upgrading the NSX Global Manager nodes. <ul style="list-style-type: none"> You must always align the version of the NSX Global Manager nodes with the rest of the SDDC stack in VMware Cloud Foundation.
VCF-LCM-REQD-004	Establish an operations practice to ensure that prior to the upgrade of any workload domain, the impact of any version upgrades is evaluated in relation to the need to upgrade NSX Global Manager.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.
VCF-LCM-REQD-005	Establish an operations practice to ensure that prior to the upgrade of the NSX Global Manager, the impact of any version change is evaluated against the existing NSX Local Manager nodes and workload domains.	The versions of NSX Global Manager and NSX Local Manager nodes must be compatible with each other. Because SDDC Manager does not provide life cycle operations (patching, update, or upgrade) for the NSX Global Manager nodes, upgrade to an unsupported version cannot be prevented.	The administrator must establish and follow an operations practice by using a runbook or automated process to ensure a fully supported and compliant bill of materials prior to any upgrade operation.

Information Security Design Elements for VMware Cloud Foundation

Use this list of requirements and recommendations for reference related to the management of access controls, certificates and accounts in a VMware Cloud Foundation environment.

For full design details, see [Information Security Design for](#) .

Table 163: Design Requirements for Account and Password Management for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-ACTMGT-REQD-SEC-001	Enable scheduled password rotation in SDDC Manager for all accounts supporting scheduled rotation.	<ul style="list-style-type: none"> Increases the security posture of your SDDC. Simplifies password management across your 	You must retrieve new passwords by using the API if you must use accounts interactively.

Table continued on next page

Continued from previous page

Recommendation ID	Design Recommendation	Justification	Implication
		SDDC management components.	
VCF-ACTMGT-REQD-SEC-003	Establish operational practice to rotate passwords using SDDC Manager on components that do not support scheduled rotation in SDDC Manager.	Rotates passwords and automatically remediates SDDC Manager databases for those user accounts.	None.
VCF-ACTMGT-REQD-SEC-003	Establish operational practice to manually rotate passwords on components that cannot be rotated by SDDC Manager.	Maintains password policies across components not handled by SDDC Manager password management.	None.

Table 164: Certificate Management Design Recommendations for VMware Cloud Foundation

Recommendation ID	Design Recommendation	Justification	Implication
VCF-SDDC-RCMD-SEC-001	Replace the default VMCA or signed certificates on all management virtual appliances with a certificate that is signed by an internal certificate authority.	Ensures that the communication to all management components is secure.	Replacing the default certificates with trusted CA-signed certificates from a certificate authority might increase the deployment preparation time because you must generate and submit certificate requests.
VCF-SDDC-RCMD-SEC-002	Use a SHA-2 algorithm or higher for signed certificates.	The SHA-1 algorithm is considered less secure and has been deprecated.	Not all certificate authorities support SHA-2 or higher.
VCF-SDDC-RCMD-SEC-003	Perform SSL certificate life cycle management for all management appliances by using SDDC Manager or SDDC Manager Plugin in vCenter.	SDDC Manager supports automated SSL certificate lifecycle management rather than requiring a series of manual steps.	Certificate management for NSX Global Manager instances must be done manually.

Planning and Preparation Workbook

Use this workbook as a dynamic VMware Cloud Foundation environment specification, calculating size, updating requirements according your design, and more.

The *Planning and Preparation Workbook* is a Microsoft Excel workbook that helps you gather the inputs required for deploying VMware Cloud Foundation (known as bring-up), VI workload domains, vSphere Supervisor, and VMware Aria Suite. It also provides guidance on the requirements for additional components that you can add to your VMware Cloud Foundation environment, such as VMware Aria Operations for Logs, VMware Aria Operations, VMware Aria Automation, and Workspace ONE Access.

Download the Planning and Preparation Workbook

[VMware Cloud Foundation Planning and Preparation Workbook](#)

Send Us Your Feedback on the Planning and Preparation Workbook

Share with us your experience with using the *VMware Cloud Foundation Planning and Preparation Workbook* to deploy an SDDC by following the VMware Cloud Foundation documentation.

This [VMware Cloud Foundation Planning and Preparation Workbook](#) instance is part of both the VMware Cloud Foundation 5.2 documentation and the VMware Validated Solutions guidance that are available on the [VMware Cloud Foundation Documentation](#) page. You use it as a planning checklist and to map the entries in the VMware Validated Solutions documentation to your SDDC deployment.

Getting Started with VMware Cloud Foundation

A high-level overview of the VMware Cloud Foundation™ product (also called VCF).

Intended Audience

The information in *Getting Started with VMware Cloud Foundation* is intended for data center cloud architects and cloud administrators who are familiar with:

- Concepts of virtualization and software-defined data centers (SDDCs)
- Networking and concepts such as uplinks, NICs, and IP networks
- Hardware components such as top-of-rack (ToR) switches, inter-rack switches, servers with direct attached storage, cables, and power supplies
- Methods for setting up physical racks in a data center
- Using VMware vSphere® to work with virtual machines

Related VMware Cloud Foundation Publications

The [VMware Cloud Foundation 5.2 Release Notes](#) lists the software components, new features, compatibility, and known issues in this VMware Cloud Foundation release.

You can open these documents from the [VMware Documentation](#) main page:

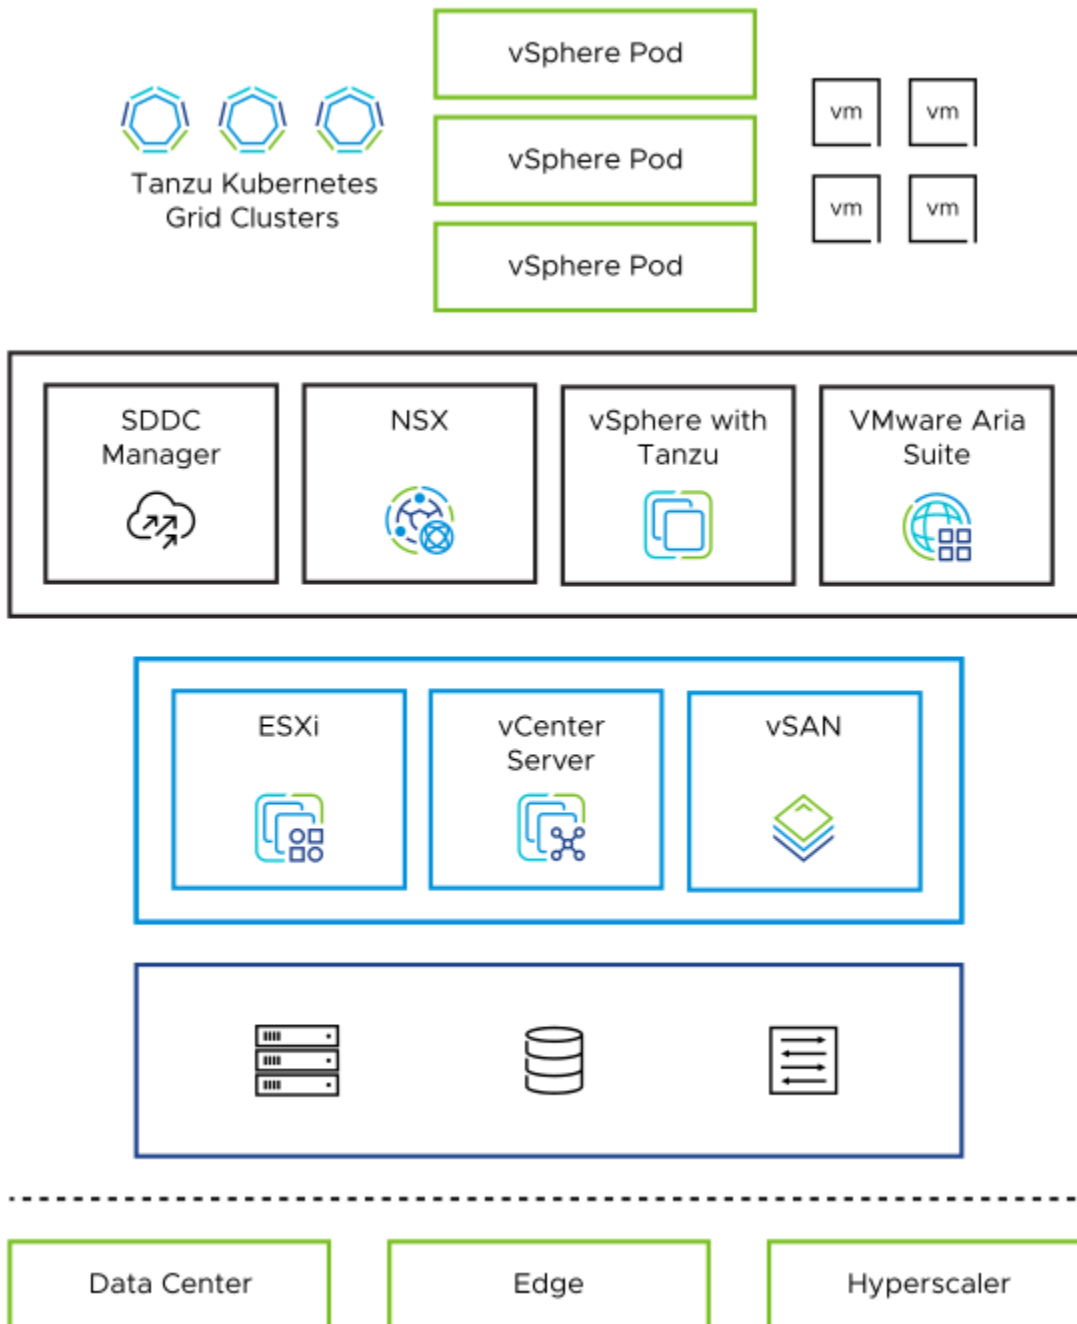
- The *VMware Cloud Foundation Planning and Preparation Workbook* provides detailed information about the inputs that are required to complete a VMware Cloud Foundation deployment. It also provides dynamic sizing guidance.
- The *VMware Cloud Foundation Deployment Guide* is intended for data center cloud administrators who deploy a VMware Cloud Foundation system in their organization's data center.
- The *VMware Cloud Foundation Administration Guide* contains detailed information about how to administer a VMware Cloud Foundation system in your data center.
- The *VMware Cloud Foundation Operations Guide* provides best practices and step-by-step instructions about certain operations in VMware Cloud Foundation, such as, full-stack shutdown and startup.
- The *VMware Cloud Foundation Lifecycle Management* document describes how to manage the life cycle of a *VMware Cloud Foundation* environment.

VMware Cloud Foundation Glossary

The [VMware Cloud Foundation Glossary](#) defines terms specific to VMware Cloud Foundation.

VMware Cloud Foundation Overview

VMware Cloud Foundation™ provides a ubiquitous hybrid cloud platform for both traditional enterprise and modern applications. Based on a proven and comprehensive software-defined stack including VMware vSphere®, VMware vSAN®, VMware NSX®, VMware vSphere® with VMware Tanzu™, and VMware Aria Suite™, VMware Cloud Foundation provides a complete set of software-defined services for compute, storage, network, container and cloud management. The result is agile, reliable, efficient cloud infrastructure that offers consistent operations across private and public clouds.



By using VMware Cloud Foundation, data center cloud administrators can provision an application environment in a rapid, repeatable, automated way versus the traditional manual process.

VMware Cloud Foundation Components

To manage the logical infrastructure in the private cloud, VMware Cloud Foundation augments the VMware virtualization and management components with VMware Cloud Builder™ and VMware Cloud Foundation™ SDDC Manager™.

VMware Cloud Foundation Component	Description
VMware Cloud Builder	VMware Cloud Builder automates the deployment of the software-defined stack, creating the first software-defined unit known as the management domain.
SDDC Manager	SDDC Manager automates the entire system life cycle, that is, from configuration and provisioning to upgrades and patching including host firmware, and simplifies day-to-day management and operations. From this interface, the virtual infrastructure administrator or cloud administrator can provision new private cloud resources, monitor changes to the logical infrastructure, and manage life cycle and other operational activities.
vSphere	<p>vSphere uses virtualization to transform individual data centers into aggregated computing infrastructures that include CPU, storage, and networking resources. VMware vSphere manages these infrastructures as a unified operating environment and provides you with the tools to administer the data centers that participate in that environment.</p> <p>The two core components of vSphere are ESXi and vCenter Server. ESXi is the virtualization platform where you create and run virtual machines and virtual appliances. vCenter Server is the service through which you manage multiple hosts connected in a network and pool host resources.</p>
vSAN	vSAN aggregates local or direct-attached data storage devices to create a single storage pool that is shared across all hosts in the vSAN cluster. Using vSAN removes the need for external shared storage, and simplifies storage configuration and virtual machine provisioning. Built-in policies allow for flexibility in data availability.
NSX	NSX is focused on providing networking, security, automation, and operational simplicity for emerging application frameworks and architectures that have heterogeneous endpoint environments and technology stacks. NSX supports cloud-native applications, bare-metal workloads, multi-hypervisor environments, public clouds, and multiple clouds.
vSphere with Tanzu	By using the integration between VMware Tanzu and VMware Cloud Foundation, you can deploy and operate the compute, networking, and storage infrastructure for vSphere IaaS Control Plane, also called Workload Management. vSphere IaaS Control Plane transforms vSphere to a platform for running Kubernetes workloads natively on the hypervisor layer. When enabled on a vSphere cluster, vSphere IaaS Control Plane provides the capability to run Kubernetes workloads directly on ESXi hosts and to create upstream Kubernetes clusters within dedicated resource pools.
VMware Aria Suite	<p>VMware Cloud Foundation supports automated deployment of VMware Aria Suite Lifecycle. You can then deploy and manage the life cycle of Workspace ONE Access and the VMware Aria Suite products (VMware Aria Operations for Logs, VMware Aria Automation, and VMware Aria Operations) by using VMware Aria Suite Lifecycle.</p> <p>VMware Aria Suite is a purpose-built management solution for the heterogeneous data center and the hybrid cloud. It is designed to deliver and manage infrastructure and applications to increase business agility while maintaining IT control. It provides the most comprehensive management stack for private and public clouds, multiple hypervisors, and physical infrastructure.</p>

For a high-level deployment process, see [Deployment Overview of](#) .

VMware Cloud Foundation Features

The VMware Cloud Foundation features provide automated deployment and life cycle management of your SDDC, and enable provisioning of customer virtualized workloads and containers.

VMware Cloud Foundation Feature	Description
Automated Software Bring-Up	You prepare your environment for VMware Cloud Foundation by installing a baseline ESXi image on vSAN ReadyNodes. After the hosts are physically racked and cabled, VMware Cloud Foundation uses the physical network details you provide (such as DNS, IP address pool, and so on) to automate the bring-up and configuration of the software stack. During bring-up, the management domain is created on the four hosts you specified. When the bring-up process completes, you have a functional management domain and can start provisioning virtual infrastructure (VI) workload domains.
Simplified Resource Provisioning with Workload Domains	In VMware Cloud Foundation, a workload domains is a policy-based resource construct with specific availability and performance attributes. See Workload Domains in .
Virtual Machines and Containers Onto the Same Platform	<p>By using the VMware Tanzu integration with VMware Cloud Foundation, you can deploy and operate the compute, networking, and storage infrastructure for vSphere with Tanzu, also called Workload Management. vSphere with Tanzu transforms vSphere to a platform for running Kubernetes workloads natively on the hypervisor layer. When enabled on a vSphere cluster, vSphere with Tanzu provides the capability to run Kubernetes workloads directly on VMware ESXi™ hosts and to create upstream Kubernetes clusters within dedicated resource pools.</p> <p>The Kubernetes concept of namespace is integrated into vSphere and becomes the unit of management. By grouping VMs and containers into logical applications via namespaces, Virtual Infrastructure (VI) admins who used to manage thousands of VMs can now manage just dozens of applications which is a massive reduction in cognitive load.</p> <p>For more information about integrating VMware Cloud Foundation with vSphere with Tanzu, see Developer Ready Infrastructure for VMware Cloud Foundation.</p>
Automated Life Cycle Management	<p>VMware Cloud Foundation offers automated life cycle management on a per-workload basis. Available updates for all components are tested for interoperability and bundled with the necessary logic for proper installation order. The update bundles are then scheduled for automatic installation on a per-workload domain basis. This allows administrators to target specific workloads or environments, for example development vs. production, for updates independent from the rest of the environment.</p> <p>vSphere Lifecycle Manager, a vCenter Server service, is integrated with VMware Cloud Foundation. By using vSphere Lifecycle Manager, you can create cluster images for centralized and simplified life cycle management of ESXi hosts including firmware. When you select the image-based life cycle management mode at VI workload domain</p>

Table continued on next page

Continued from previous page

VMware Cloud Foundation Feature	Description
	creation, you can update and upgrade the ESXi version on all hosts in the cluster collectively. You can also install and update vendor add-ons and components on all ESXi hosts in a cluster. See Managing vSphere Lifecycle Manager Images in VMware Cloud Foundation .
Stretched Deployment	You can set up two availability zones in your environment and introduce high availability of management and customer workloads by configuring vSAN stretched clusters by using the SDDC Manager API. Availability zones protect against failures of groups of hosts. These group can consist of hosts in the same data center, for example, installed in different racks, chassis or rooms, or in different data centers with low-latency high-speed links connecting them. Using two availability zones can improve availability of management components running the SDDC, minimize downtime of services, and improve SLAs. See Managing vSphere Lifecycle Manager Images in VMware Cloud Foundation .
NSX Federation	<p>You can use NSX Federation to propagate configurations that span multiple NSX instances in a single VMware Cloud Foundation instance or across multiple VMware Cloud Foundation instances. You can set up global networking, enabling failover of segment ingress and egress traffic between VMware Cloud Foundation instances, and implement a unified firewall configuration.</p> <p>In the management domain in a deployment with multiple VMware Cloud Foundation instances, you use NSX to provide cross-instance services to SDDC management components which do not have native support for availability at several locations, such as VMware Aria Automation and VMware Aria Operations. In a management domain, you can use NSX Federation only to connect to the management domains of other VMware Cloud Foundation instances. Avoid connecting a management domain with VI workload domains in a single NSX Federation instance.</p> <p>You configure NSX Federation in VMware Cloud Foundation manually.</p> <p>For more information on using NSX Federation with VMware Cloud Foundation, see NSX Design for VMware Cloud Foundation and Working with NSX Federation in VMware Cloud Foundation.</p>

VMware Cloud Foundation Glossary

In VMware Cloud Foundation, you perform specific operations and use unique constructs for automated SDDC deployment and maintenance.

Term	Description
availability zone	A collection of infrastructure components. Each availability zone is isolated from the other availability zones to prevent the propagation of failure or outage across the data center. In VMware Cloud Foundation, you implement availability of workloads across availability zones by using vSAN stretched clusters.

Table continued on next page

Continued from previous page

Term	Description
Application virtual networks (AVNs)	Virtual networks backed by overlay or VLAN NSX segments using the encapsulation protocol of VMware NSX. An AVN uses a single IP address space to span across data centers.
bring-up	Deployment and initial configuration of a VMware Cloud Foundation system. During the bring-up process, the management domain is created and the VMware Cloud Foundation software stack is deployed on the management domain.
commission a host	Adding a host to VMware Cloud Foundation inventory. The host becomes unassigned.
dirty host	A host that has been removed from a cluster in a workload domain. A dirty host cannot be assigned to another workload domain until it is decommissioned, re-imaged, and commissioned again.
decommission a host	Removing an unassigned host from the VMware Cloud Foundation inventory. SDDC Manager does not manage decommissioned hosts.
NSX Edge cluster	A logical grouping of NSX Edge nodes. These nodes run on a vSphere cluster, and provide north-south and east-west routing and network services for the management or VI workload domain.
free pool	Hosts in the VMware Cloud Foundation inventory that are not assigned to a workload domain.
host	A server that is imaged with the ESXi software.
install bundle	Contains software for VI workload domains and VMware Aria Suite Lifecycle. You can use an install bundle to deploy later versions of the software components in a new VI workload domain than the versions in the Bill of Materials for VMware Cloud Foundation.
inventory	Logical and physical entities managed by VMware Cloud Foundation.
Kubernetes - Workload Management	With Kubernetes - Workload Management, you can deploy and operate the compute, networking, and storage infrastructure for vSphere IaaS Control Plane workloads. A vSphere IaaS Control Plane workload is an application with containers running inside vSphere pods, regular VMs, or Tanzu Kubernetes clusters.
Lifecycle Manager (LCM)	Automates patching and upgrading of the software stack.
management domain	One or more vSphere clusters of physical hosts that contain the management component VMs, such as vCenter Server, NSX Manager cluster, management NSX Edge cluster, SDDC Manager, and so on. The management domain supports only vSAN storage.
network pool	Automatically assigns static IP addresses to vSAN and vMotion VMkernel ports so that you don't need to enter IP addresses manually when creating a VI workload domain or adding a host or cluster to a workload domain.
update bundle	Contains software to update the VMware Cloud Foundation components in your management or VI workload domain.
principal storage	Required for each vSphere cluster, containing the data of the virtual machines in the cluster. For the management domain, only vSAN principal storage is supported. For a VI workload domain, you set the principal storage when creating the domain or when adding a cluster to the domain. You cannot change the principal storage later. See also <i>supplemental storage</i> .
SDDC Manager	A software component that provisions, manages, and monitors the logical and physical resources of a VMware Cloud Foundation system. SDDC Manager provides the user interface for managing VMware Cloud Foundation, CLI-based administrator tools, and an API for further automation.
server	A bare-metal server in a physical rack. After imaging, it is referred to as a host.

Table continued on next page

Continued from previous page

Term	Description
supplemental storage	Extends the capacity of the workload domain for hosting more virtual machines or storing supporting data, such as backups. You can add or remove supplemental storage to clusters in the management or VI workload domain at any time.
unassigned host	A host in the free pool that does not belong to a workload domain.
vSphere Lifecycle Manager (vLCM)	A vCenter Server service, which is integrated with VMware Cloud Foundation, that enables centralized and simplified life cycle management of ESXi hosts.
virtual infrastructure (VI) workload domain	One or more vSphere clusters that contain customer workloads. VMware Cloud Foundation scales and manages the life cycle of each VI workload domain independently. The vCenter Server instance and NSX Manager cluster for a VI workload domain are physically located in the management domain, while the NSX edge nodes - on the VI workload domain.
vSphere Lifecycle Manager baseline	A grouping of multiple bulletins. You can attach a baseline to an ESXi host and check the compliance of the host against the associated baseline. According to the type of content, baselines are patch baselines, extension baselines, and upgrade baselines. SDDC Manager creates the required baseline and baseline group for updating a cluster in a workload domain.
vSphere Lifecycle Manager image	A precise description of the software, components, vendor add-ons, and firmware to run on an ESXi host. You set up a single image and apply it to all hosts in a cluster, thus ensuring cluster-wide host image homogeneity.
workload domain	<p>A policy-based resource container with specific availability and performance attributes that combines vSphere, storage (vSAN, NFS, VMFS on FC, or vVols) and networking (VMware NSX) into a single consumable entity. A workload domain can be created, expanded, and deleted as part of the SDDC life cycle operations. It can contain clusters of physical hosts with a corresponding vCenter Server instance to manage them.</p> <p>VMware Cloud Foundation supports two types of workload domains - the management domain and one or more VI workload domains.</p>

Deployment Guide

Provides information about installing VMware ESXi software on VMware Cloud Foundation servers and deploying the management domain using the VMware Cloud Builder appliance.

About the VMware Cloud Foundation Deployment Guide

The *VMware Cloud Foundation Deployment Guide* provides information about installing VMware ESXi™ software on VMware Cloud Foundation™ servers and deploying the management domain using the VMware Cloud Builder appliance™. Starting with VMware Cloud Foundation 5.2, you can also use the VCF Import Tool to convert an existing vSphere cluster into a management domain.

Intended Audience

The *VMware Cloud Foundation Deployment Guide* is intended for data center cloud administrators who deploy a VMware Cloud Foundation system in their organization's data center. The information in this guide is written for experienced data center cloud administrators who are familiar with:

- Concepts of virtualization and software-defined data centers
- Networking and concepts such as uplinks, NICs, and IP networks
- Hardware components such as top-of-rack (ToR) switches, inter-rack switches, servers with direct attached storage, cables, and power supplies
- Methods for setting up physical racks in your data center
- Using the VMware vSphere® Client™ to work with virtual machines

Related Publications

Getting Started with VMware Cloud Foundation document provides a high-level overview of the product

The *Planning and Preparation Workbook* provides detailed information about the software, tools, and external services that are required for Cloud Foundation.

The *VMware Cloud Foundation Administration Guide* contains detailed information about how to administer and operate a VMware Cloud Foundation system in your data center. It also contains information about using the VCF Import Tool to convert an existing vSphere environment to the VMware Cloud Foundation management domain.

Your VMware Cloud Foundation system includes various VMware software products and components. You can find the documentation for those VMware software products at docs.vmware.com.

VMware Cloud Foundation Glossary

The VMware Cloud Foundation Glossary defines terms specific to VMware Cloud Foundation.

Preparing your Environment for VMware Cloud Foundation

Before you start the automated deployment of the management domain using VMware Cloud Builder, your environment must meet target prerequisites and be in a specific starting state.

Prepare the platform by deploying and configuring the necessary infrastructure components. For detailed prerequisites, see the *Planning and Preparation Workbook*.

Deploying VMware Cloud Foundation

You begin the VMware Cloud Foundation deployment process by deploying the VMware Cloud Builder appliance. After imaging your servers, you download and complete the deployment parameters workbook from the VMware Cloud Builder appliance to define your network information, host details, and other required information. During the deployment process,

this workbook is uploaded to the VMware Cloud Builder appliance, where a JSON file is generated to drive the bring-up process. The provided information is validated, and the automated phase of the bring-up process begins.

You must prepare your environment for deploying VMware Cloud Foundation. See the *Planning and Preparation Workbook*.

You can perform bring-up with certificates generated by an external CA, in which case ESXi certificates are not replaced with vCenter Server signed certificates. If you use external certificates for ESXi hosts in the management domain, hosts added after bring-up must also be added with external certificates. This feature is supported only through APIs. For more information, see [Deploy the Management Domain Using ESXi Hosts with External Certificates](#).

Deploy VMware Cloud Builder Appliance

VMware Cloud Builder is a virtual appliance that is used to deploy and configure the first cluster of the management domain and transfer inventory and control to SDDC Manager. During the deployment process, the VMware Cloud Builder appliance validates network information you provide in the deployment parameter workbook such as DNS, network (VLANS, IPs, MTUs), and credentials.

Before you deploy the VMware Cloud Builder appliance, verify that your environment fulfills the requirements for this process.

Prerequisite	Value
Environment	<ul style="list-style-type: none"> Verify that your environment is configured for deployment of VMware Cloud Builder and the management domain. Verify that you have available virtual infrastructure that has access to the management network that will be used by the management domain. You deploy VMware Cloud Builder on that virtual infrastructure.
Resource Requirements	<ul style="list-style-type: none"> 4 CPUs 4 GB of Memory 279 GB of Storage <ul style="list-style-type: none"> – 25.1 GB (thin provisioned) – 253.8 GB (thick provisioned)
Installation Packages	Verify that you download the OVA file(s) for VMware Cloud Builder.
Network	<ul style="list-style-type: none"> Verify that the static IP address and FQDN for the VMware Cloud Builder appliance are available. Verify that connectivity is in place from the VMware Cloud Builder appliance and the management VLAN used in the deployment.

To automate the deployment, the VMware Cloud Builder appliance must be on the same management network as the hosts to be used. It must also be able to access all required external services, such as DNS and NTP.

You must deploy the VMware Cloud Builder appliance on a suitable platform. This can be on a laptop running VMware Workstation or VMware Fusion, or on an ESXi host. The VMware Cloud Builder appliance must have network access to all hosts on the management network.

This procedure describes how to deploy the VMware Cloud Builder appliance directly to an ESXi host.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the navigation pane, select **Host**, and click **Create/Register VM**.
3. On the Select creation type dialog box, select **Deploy a virtual machine from an OVF or OVA file** and click **Next**.

4. On the Select OVF and VMDK files page, enter a name for the virtual machine, select the VMware Cloud Builder.ova file, and click **Next**.
5. On the Select Storage page, select a datastore and click **Next**.
6. On the License agreements dialog box, click **I agree** and then click **Next**.
7. On the Select networks dialog box, enter the following values and click **Next**.

Setting	Value
Network mappings	<i>your_portgroup</i>
Disk provisioning	Thin
Power on automatically	Selected

8. On the Additional settings dialog box, expand **Application**, enter the following values, and click **Next**.

Setting	Details
Admin Username	Accept the default admin user name, <code>admin</code> .
Admin Password/Admin Password confirm	The admin password must be a minimum of 15 characters and include at least one uppercase, one lowercase, one digit, and one special character. Supported special characters: @ ! # \$ % ? ^ NOTE A password cannot be based on a dictionary word (for example, <code>VMware1!</code>)
Root password/Root password confirm	The root password must be a minimum of 15 characters and include at least one uppercase, one lowercase, one digit, and one special character. Supported special characters: @ ! # \$ % ? ^ NOTE A password cannot be based on a dictionary word (for example, <code>VMware1!</code>)
Hostname	Enter the hostname for the VMware Cloud Builder appliance.
Network 1 IP Address	Enter the IP address for the VMware Cloud Builder appliance.
Network 1 Subnet Mask	Enter the subnet mask for the VMware Cloud Builder appliance.
Default Gateway	Enter the default gateway for the VMware Cloud Builder appliance.
DNS Servers	Enter the IP address of the primary and secondary DNS servers (comma separated). Do not specify more than two servers.
DNS Domain Name	Enter the DNS domain name. For example, <code>vsphere.local</code> .
DNS Domain Search Paths	Enter the DNS domain search path(s). Use a comma if entering multiple search paths. For example <code>vsphere.local, sfo.vsphere.local</code> .

Table continued on next page

Continued from previous page

Setting	Details
NTP Servers	Enter the NTP server(s). Use a comma if entering multiple NTP servers. NTP servers can be entered using FQDNs or IP addresses.

9. On the Ready to complete page, review the virtual machine configuration and click **Finish**.

NOTE

Make sure your passwords meet the requirements specified above before clicking **Finish** or your deployment will not succeed.

10. After the VMware Cloud Builder appliance is deployed, SSH in to the VM with the admin credentials provided in step 8.
11. Ensure that you can ping the ESXi hosts.
12. Verify that the VMware Cloud Builder appliance has access to the required external services, such as DNS and NTP by performing forward and reverse DNS lookups for each host and the specified NTP servers.

Prepare ESXi Hosts for VMware Cloud Foundation

Before you can begin the process of deploying VMware Cloud Foundation you must prepare the ESXi hosts that will form the management domain.

The management domain requires a minimum of four ESXi hosts.

To use vSAN Express Storage Architecture (ESA), your hosts must be ESA-compatible.

TIP

See the [vSAN ESA VCG](#) for information about compatible hardware.

Preparing the ESXi hosts involves installing the correct version of ESXi and performing some basic configuration tasks. For the supported ESXi version, see the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes*.

Create a Custom ISO Image for ESXi

When your environment requires a custom ISO file for ESXi, you can create one using VMware PowerCLI or vSphere Lifecycle Manager.

Download the zip files for the following:

- ESXi patch for the ESXi version specified in the VMware Cloud Foundation BOM or in the list of supported async patches in [KB 88287](#). You can download patches from the Broadcom Support Portal.

NOTE

If you are preparing hosts for a VI workload domain where the ESXi hosts have been async patched to a later version of ESXi than the version listed in the BOM, the new hosts must use the later version of ESXi.

- OEM add-on for ESXi from the Broadcom Support Portal. If the ESXi version specified in the BOM is not available in the **Select Version** drop-down menu, contact your vendor to determine which OEM add-on version to use.

You might need to create a custom ISO image for ESXi in the following situations:

- The ESXi version specified in the VMware Cloud Foundation BOM does not have an associated ISO file on the Broadcom Support Portal. This can be the case for ESXi patch releases.
- You need an async patch version of ESXi.
- You need a vendor-specific (OEM) ISO file.

Create a Custom ESXi ISO Image Using VMware PowerCLI

You can use VMware Power CLI to create a custom ISO.

VMware PowerCLI 12.0 or later.

1. Gather the required information for the software spec that is used to create the custom ISO.
 - a) In VMware PowerCLI, use the [Get-DepotBaseImages](#) cmdlet to get the base image version from the zip file for the ESXi patch that you downloaded from the patches portal.

For example:

```
Get-DepotBaseImages "c:\temp\VMware-ESXi-7.0U1d-17551050-depot.zip"
```

- b) Use the [Get-DepotAddons](#) cmdlet to get the add-on name and version from the zip file for the OEM add-on for ESXi that you downloaded from the Broadcom Support Portal. (if applicable)

For example:

```
Get-DepotAddons "c:\temp\HPE-701.0.0.10.6.5.12-Jan2021-Synergy-Addon-depot.zip"
```

2. Create the software spec using the information you gathered in step 1.

The software spec is a JSON file that contains information about the ESXi version and vendor add-on (if applicable). For example:

```
{
  "add_on": {
    "name": "HPE-Custom-Syn-AddOn",
    "version": "701.0.0.10.6.5-12"
  },
  "base_image": {
    "version": "7.0.1-0.30.17551050"
  },
  "components": null,
  "hardware_support": null,
  "solutions": null
}
```

3. In VMware PowerCLI, use the `New-IsoImage` cmdlet to generate a custom ISO.

For example:

```
New-IsoImage -SoftwareSpec "c:\temp\HPE-70U1d-custom.JSON" -Depots "c:\temp\VMware-ESXi-7.0U1d-17551050-depot.zip" , "c:\temp\HPE-701.0.0.10.6.5.12-Jan2021-Synergy-Addon-depot.zip" -Destination "c:\temp\HPE-70U1d-custom.iso"
```

Provide the path to the software spec you created in step 2.

The depot(s) include the path to the zip files for the supported ESXi version and vendor add-on.

The destination include the path and file name for the custom ISO file.

For more information about the `New-IsoImage` cmdlet, see <https://developer.broadcom.com/powercli/latest/vmware.imagebuilder/commands/new-isoimage>.

Create a Custom ESXi ISO Image Using vSphere Lifecycle Manager

If you have access to a vCenter Server 7.0 environment, you can use vSphere Lifecycle Manager to create and export a custom ISO.

Import the ESXi patch and vendor add-on (if applicable) zip files to the vSphere Lifecycle Manager depot. See [Import Updates to the vSphere Lifecycle Manager Depot](#).

1. Log in to vCenter Server using the vSphere Client.
2. Create a new temporary cluster, selecting the **Manage all hosts in the cluster with a single image** check box.
3. Select the ESXi version and vendor add-on (optional) and click **OK**.
4. Export the vSphere Lifecycle Manager image as an ISO.
See [Export an Image](#).
5. Delete the temporary cluster.

Install ESXi Interactively and Configure Hosts for VMware Cloud Foundation

You can interactively install ESXi on all the hosts that will form the first cluster in the management domain, then you configure the management network, DNS, and NTP services. You can use the same process to add more hosts to the management domain later, or to install and configure hosts for VI workload domains.

- Download the ESXi ISO from the Broadcom Support Portal. For the supported ESXi versions, see the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes* and the list of supported async patches in [KB 88287](#). If the required version of ESXi does not have an ISO available on the Broadcom Support Portal, you can create one. See [Create a Custom ISO Image for ESXi](#).

NOTE

If you are preparing hosts for a VI workload domain where the ESXi hosts have been async patched to a later version of ESXi than the version listed in the BOM, the new hosts must use the later version of ESXi.

- Make sure that you have a host machine for SDDC access. You use this host to connect to the data center and perform configuration steps.
- Verify that you have the completed *Planning and Preparation Workbook*.
- Verify the Prerequisite Checklist sheet in the *Planning and Preparation Workbook*.

ESXi 8.0 Update 3 and later support installing two data processing units (DPUs) for use with VMware Cloud Foundation 5.2 or later.

You can utilize the two DPUs in Active/Standby mode to provide high availability. Such configuration provides redundancy in the event one of the DPUs fails. In the high availability configuration, both DPUs are assigned to the same NSX-backed vSphere Distributed Switch. For example, DPU-1 is attached to vmnic0 and vmnic1 of the vSphere Distributed Switch and DPU-2 is attached to vmnic2 and vmnic3 of the same vSphere Distributed Switch.

You can also utilize the two DPUs as independent devices to increase offload capacity per ESXi host. Each DPU is attached to a separate vSphere Distributed Switch and you have no failover between DPUs in such configuration.

Install ESXi on VMware Cloud Foundation Hosts Using the ISO

Install ESXi on all hosts in the first cluster in the management domain interactively. You can use the same process to install ESXi on additional hosts for the management domain, or on hosts for a VI workload domain.

Repeat this procedure for all hosts in the first cluster in the management domain.

1. Mount the ESXi ISO on the host and restart the machine.
2. Set the BIOS or UEFI to boot from the mounted ISO.

NOTE

If your system has supported data processing units (DPUs), you can only use UEFI to install and boot ESXi on the DPUs.

See your hardware vendor documentation for information on changing boot order.

3. On the welcome screen, press **Enter** to continue.
4. Accept the End User License Agreement by pressing **Enter**.

Starting with ESXi 8.0 Update 3, after the scanning for available devices completes, if your system has DPUs, you see them automatically listed with their respective PCI slots. You no longer select a slot. The DPU devices must be identical: same vendor, same hardware version and same firmware

5. On the **Select a Disk to Install or Upgrade** screen, select the drive on which to install ESXi on and press **Enter**.
6. Select the keyboard type for the host.

You can change the keyboard type after installation in the direct console.

7. Enter the root password for the host.
8. In the **Confirm Install** screen, if you have DPUs, you see each listed on a separate row. Press **F11** to confirm the start of the installation.

Starting with ESXi 8.0 Update 3, if your systems has DPUs, you see a single progress bar for the ESXi and DPU installation, with dynamic updates to the label showing what stage of the installer is being run.

9. On the **Installation Complete** screen, press **Enter** to reboot the host.
10. Set the first boot device to be the drive on which you installed ESXi.
11. Repeat this procedure for all remaining hosts.

Configure the Network on VMware Cloud Foundation Hosts

After the initial boot, use the ESXi Direct Console User Interface (DCUI) for host network configuration and administrative access.

Perform the following tasks to configure the host network settings:

- Configure the network adapter (vmk0) and VLAN ID for the Management Network.
- Configure the IP address, subnet mask, gateway, DNS server, and FQDN for the ESXi host.

Repeat this procedure for all hosts that you are adding to the first cluster of the management domain. Enter the respective values from the completed *Planning and Preparation Workbook*.

1. Open the DCUI of the ESXi host.
 - a) Open a console window to the host.
 - b) Press F2 to enter the DCUI.
 - c) Log in by using the `esxi_root_user_password`.
2. Configure the network.
 - a) Select **Configure Management Network** and press Enter.
 - b) Select **VLAN (Optional)** and press Enter.
 - c) Enter the VLAN ID for the Management Network and press Enter.
 - d) Select **IPv4 Configuration** and press Enter.
 - e) Select **Set static IPv4 address and network configuration** and press the Space bar.

- f) Enter the IPv4 Address, Subnet Mask and Default Gateway and press Enter.
 - g) Select **DNS Configuration** and press Enter.
 - h) Select **Use the following DNS Server address and hostname** and press the Space bar.
 - i) Enter the Primary DNS Server, Alternate DNS Server and Hostname (FQDN) and press Enter.
 - j) Select **Custom DNS Suffixes** and press Enter.
 - k) Ensure that there are no suffixes listed and press Enter.
3. Press Escape to exit and press Y to confirm the changes.
 4. Repeat this procedure for all remaining hosts.

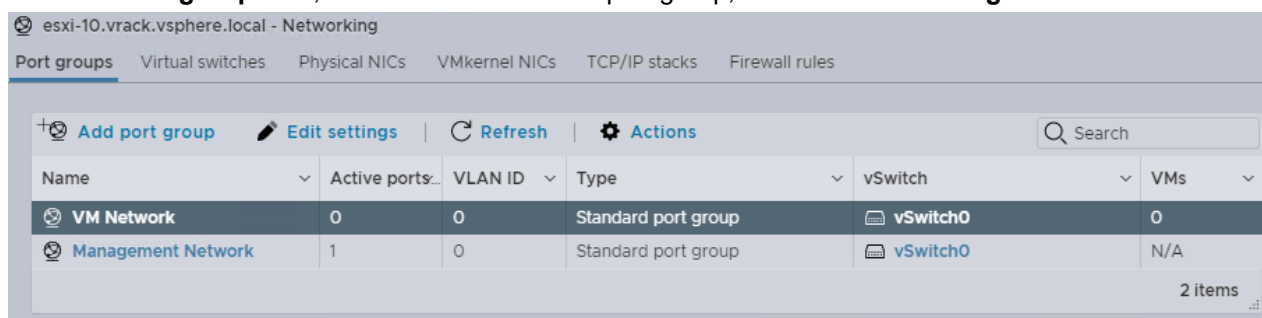
Configure the Virtual Machine Network Port Group on VMware Cloud Foundation Hosts

You perform configuration of the Virtual Machine Network port group for each ESXi host by using the VMware Host Client.

You configure the VLAN ID of the VM Network port group on the vSphere Standard Switch. This configuration provides connectivity to the Management network to allow communication to the vCenter Server Appliance during the automated deployment.

Repeat this procedure for all hosts in the first cluster of the management domain. Enter the respective values from the completed *Planning and Preparation Workbook*.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. Click **OK** to join the Customer Experience Improvement Program.
3. Configure a VLAN for the VM Network port group.
 - a) In the navigation pane, click **Networking**.
 - b) Click the **Port groups** tab, select the **VM network** port group, and click **Edit Settings**.



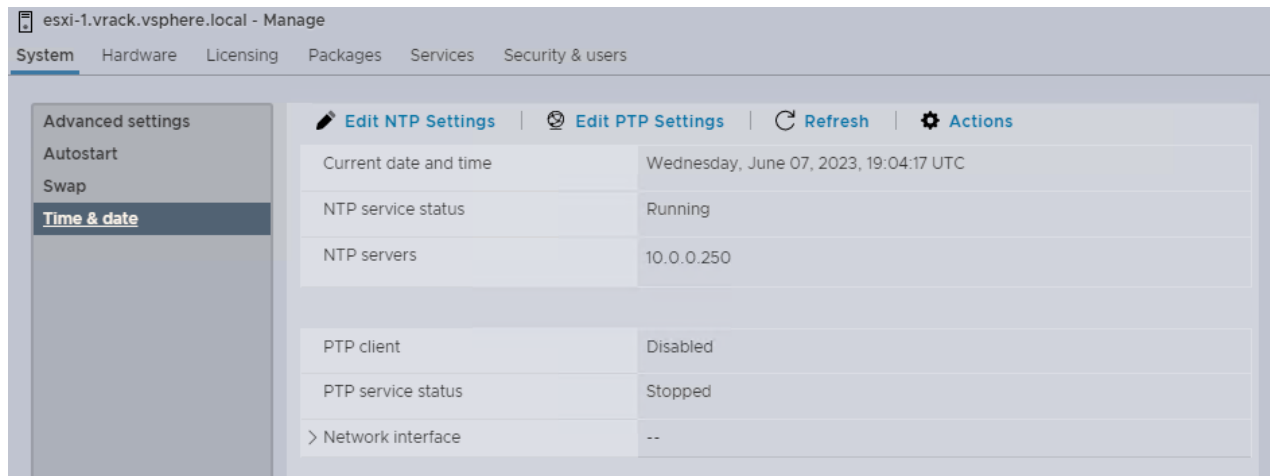
- c) On the **Edit port group - VM network** page, enter the Management Network VLAN ID, and click **Save**.
4. Repeat this procedure for all remaining hosts.

Configure NTP on VMware Cloud Foundation Hosts

Complete the initial configuration of all ESXi hosts by configuring the NTP service to avoid time synchronization issues in the SDDC.

Repeat this procedure for all hosts in the first cluster of the management domain. Enter the respective values from the completed *Planning and Preparation Workbook*.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. Configure and start the NTP service.
 - a) In the navigation pane, click **Manage**, and click the **System** tab.



- b) Click **Time & date** and click **Edit NTP Settings**.
 - c) On the **Edit NTP Settings** page, select the **Use Network Time Protocol (enable NTP client)** radio button, and change the NTP service startup policy to **Start and stop with host**.
 - d) In the **NTP servers** text box, enter the NTP Server FQDN or IP Address, and click **Save**.
 - e) To start the service, click **Actions**, select **NTP service**, and click **Start**.
3. Repeat this procedure for all remaining hosts.

Regenerate the Self-Signed Certificate on All Hosts

Once you have configured the ESXi hosts' identity by providing a hostname you must regenerate the self-signed certificate to ensure the correct common name is defined.

During the installation of ESXi, the installer generates a self-signed certificate for each ESXi host but the process is performed prior to the ESXi identity being configured. This means all ESXi hosts have a common name in their self-signed certificate of `localhost.localdomain`. All communication between VMware Cloud Builder and the ESXi hosts is performed securely over HTTPS and as a result it validates the identity when making a connection by comparing the common name of the certificate against the FQDN provided within the VMware Cloud Builder configuration file. To ensure that the connection attempts and validation does not fail, you must manually regenerate the self-signed certificate after hostname has been configured.

NOTE

VMware Cloud Foundation supports the use of signed certificates. If your organization's security policy mandates that all ESXi hosts must be configured with a CA-signed certificate, see [Configure ESXi Hosts with Signed Certificates](#).

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the **Actions** menu, click **Services > Enable Secure Shell (SSH)**.
3. Log in to the ESXi host using an SSH client such as Putty.
4. Regenerate the self-signed certificate by executing the following command:

```
/sbin/generate-certificates
```

5. Reboot the host.
6. Log back in to the VMware Host Client and click **Services > Disable Secure Shell (SSH)** from the **Actions** menu.
7. Repeat this procedure for all remaining hosts.

Configure ESXi Hosts with Signed Certificates

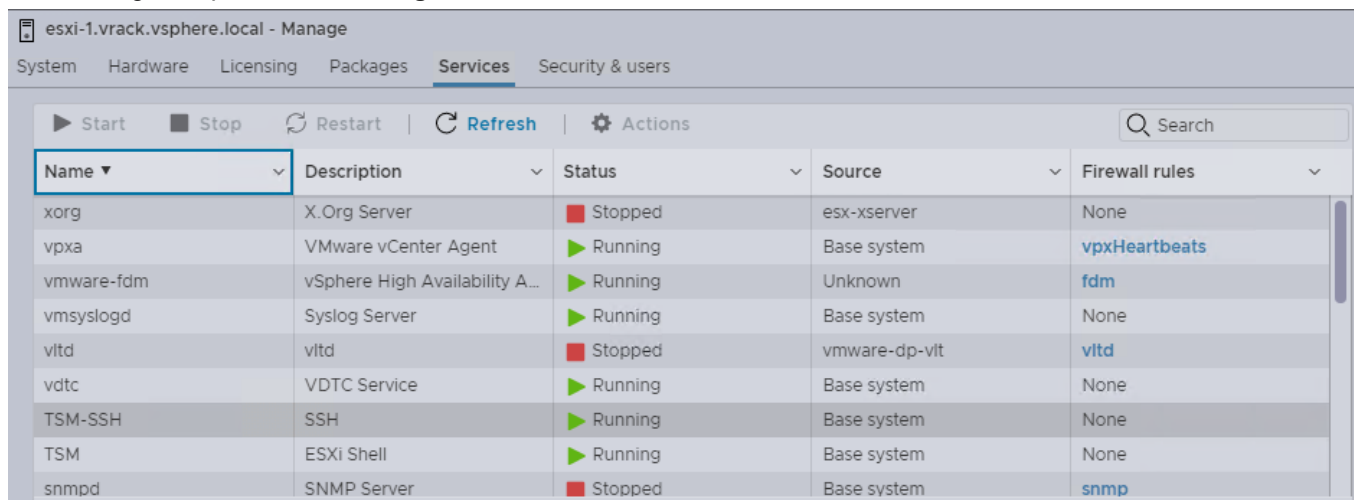
If corporate policy requires that you use external CA-signed certificates instead of VMCA-signed certificates for ESXi hosts, you can manually add external certificates to the hosts.

External CA-signed certificate and key are available.

When you install ESXi software on a server to create an ESXi host, the host initially has an autogenerated certificate. By default, when the host is added to a vCenter Server system during bring-up of the management domain or other operations involving hosts (for example, host commissioning, VI workload domain creation, and so on), the autogenerated certificate is replaced with a certificate that is signed by the VMware Certificate Authority (VMCA).

When you use external certificates during bring-up, they are not replaced by VMCA-signed certificates. Once you perform bring-up with external certificates for ESXi hosts, all future hosts added to VMware Cloud Foundation must also use external certificates.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the navigation pane, click **Manage** and click the **Services** tab.



The screenshot shows the VMware Host Client interface for an ESXi host. The 'Services' tab is selected, displaying a table of services. The table has columns for Name, Description, Status, Source, and Firewall rules. The 'TSM-SSH' service is highlighted in the list.

Name	Description	Status	Source	Firewall rules
xorg	X.Org Server	Stopped	esx-xserver	None
vpax	VMware vCenter Agent	Running	Base system	vpHeartbeats
vmware-fdm	vSphere High Availability A...	Running	Unknown	fdm
vmsyslogd	Syslog Server	Running	Base system	None
vitd	vitd	Stopped	vmware-dp-vit	vitd
vdtc	VDTC Service	Running	Base system	None
TSM-SSH	SSH	Running	Base system	None
TSM	ESXi Shell	Running	Base system	None
snmpd	SNMP Server	Stopped	Base system	snmp

3. Select the **TSM-SSH** service and click **Start** if not started.
4. Log in to the ESXi Shell for the first host, either directly from the DCUI or from an SSH client, as a user with administrator privileges.
5. In the directory `/etc/vmware/ssl`, rename the existing certificates using the following commands:


```
mv rui.crt orig.rui.crt
mv rui.key orig.rui.key
```
6. Copy the external certificate and key that you want to use to `/etc/vmware/ssl`.
7. Rename the external certificate and key to `rui.crt` and `rui.key`.
8. Restart the host management agents by running the following commands:


```
/etc/init.d/hostd restart
/etc/init.d/vpax restart
```
9. In the VMware Host Client, select the **TSM-SSH** service for the ESXi host and click **Stop**.
10. Repeat for all the ESXi hosts that you are adding to VMware Cloud Foundation.

See [Deploy the Management Domain Using ESXi Hosts with External Certificates](#).

Deploy the Management Domain Using VMware Cloud Builder

The VMware Cloud Foundation deployment process is referred to as bring-up. You specify deployment information specific to your environment such as networks, hosts, license keys, and other information in the deployment parameter workbook and upload the file to the VMware Cloud Builder appliance to initiate bring-up of the management domain.

During bring-up, the management domain is created on the ESXi hosts specified in the deployment parameter workbook. The VMware Cloud Foundation software components are automatically deployed, configured, and licensed using the information provided. The deployment parameter workbook can be reused to deploy multiple VMware Cloud Foundation instances of the same version.

The following procedure describes how to perform bring-up of the management domain using the deployment parameter workbook. You can also perform bring-up using a custom JSON specification. See the [VMware Cloud Foundation API Reference Guide](#) for more information.

NOTE

Starting with VMware Cloud Foundation 5.2, you can use the VCF Import Tool to convert an existing vSphere environment to create the management domain. See [Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#).

Some use cases are only available using a custom JSON specification. For example, using custom CA-signed certificates for ESXi hosts. See [Deploy the Management Domain Using ESXi Hosts with External Certificates](#).

1. In a web browser, log in to the VMware Cloud Builder appliance administration interface: `https://Cloud_Builder_VM_FQDN`.
2. Enter the admin credentials you provided when you deployed the VMware Cloud Builder appliance and then click **Log In**.
3. On the **End-User License Agreement** page, select the **I Agree to the End User License Agreement** check box and click **Next**.
4. Select **VMware Cloud Foundation** and click **Next**.
5. Review and acknowledge the prerequisites and click **Next**.

If there are any gaps, ensure they are fixed before proceeding to avoid issues during the bring-up process. You can download or print the prerequisite list for reference.

6. Download the deployment parameter workbook from the [Broadcom Support portal](#) and fill it in with the required information.

See [About the Deployment Parameter Workbook](#).

7. Click **Next**.
8. Click **Select File**, browse to the completed workbook, and click **Open** to upload the workbook.
9. Click **Next** to begin validation of the uploaded file.

To access the bring-up log file, SSH to the VMware Cloud Builder appliance as `admin` and open the `/opt/vmware/bringup/logs/vcf-bringup-debug.log` file.

If there is an error during the validation and the **Next** button is grayed out, you can either make corrections to the environment or edit the deployment parameter workbook and upload it again. Then click **Retry** to perform the validation again.

If any warnings are displayed and you want to proceed, click **Acknowledge** and then click **Next**.

10. Click **Deploy SDDC**.

During the bring-up process, the vCenter Server, NSX, and SDDC Manager appliances are deployed and the management domain is created. The status of the bring-up tasks is displayed in the UI.

After bring-up is completed, a green bar is displayed indicating that bring-up was successful. A link to the SDDC Manager UI is also displayed. If there are errors during bring-up, see [Troubleshooting Deployment](#).

11. Click **Download** to download a detailed deployment report. This report includes information on assigned IP addresses and networks that were configured in your environment.
12. After bring-up is completed, click **Finish**.
13. Click **Launch SDDC Manager**.
14. Power off the VMware Cloud Builder appliance.

About the Deployment Parameter Workbook

The deployment parameter workbook contains worksheets categorizing the information required for deploying VMware Cloud Foundation. The information provided is used to create the management domain using the VMware Cloud Builder appliance.

Before you begin filling in the deployment parameter workbook, download the workbook from the [Broadcom Support portal](#).

The fields in yellow contain sample values that you should replace with the information for your environment. If a cell turns red, the required information is missing, or validation input has failed.

IMPORTANT

The deployment parameter workbook is not able to fully validate all inputs due to formula limitations of Microsoft Excel. Some validation issues may not be reported until you upload the deployment parameter workbook to the VMware Cloud Builder appliance.

NOTE

Do not copy and paste content between cells in the deployment parameter workbook, since this may cause issues.

The Introduction worksheet in the deployment parameter workbook contains an overview of the workbook and guidance on how to complete it. For information about the prerequisites for deploying the management domain, see the *Planning and Preparation Workbook*.

Credentials Worksheet

The Credentials worksheet details the accounts and initial passwords for the VMware Cloud Foundation components. You must provide input for each yellow box. A red cell may indicate that validations on the password length has failed.

Input Required

Update the Default Password field for each user (including the automation user in the last row). Passwords can be different per user or common across multiple users. The tables below provide details on password requirements.

Table 165: Password Complexity

Password	Requirements
ESXi Host root account	This is the password which you configured on the hosts during ESXi installation.
Default Single-Sign on domain administrator user	<ol style="list-style-type: none"> 1. Length 8-20 characters 2. Must include: <ul style="list-style-type: none"> – mix of upper-case and lower-case letters – a number – a special character, such as @ ! # \$ % ^ or ? 3. Must not include * { } [] () / \ ' " ` ~ , ; : . < >

Table continued on next page

Continued from previous page

Password	Requirements
vCenter Server virtual appliance root account	<ol style="list-style-type: none"> 1. Length 8-20 characters 2. Must include: <ul style="list-style-type: none"> – mix of upper-case and lower-case letters – a number – a special character, such as @ ! # \$ % ^ or ? 3. Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
NSX virtual appliance root account	<ol style="list-style-type: none"> 1. Length 12-127 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? – at least five different characters 3. Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
NSX user interface and default CLI admin account	<ol style="list-style-type: none"> 1. Length 12-127 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? – at least five different characters 3. Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
NSX audit CLI account	<ol style="list-style-type: none"> 1. Length 12-127 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? – at least five different characters 3. Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
SDDC Manager appliance root account	<ol style="list-style-type: none"> 1. Minimum length 15 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? 3. Must not include: <ul style="list-style-type: none"> – * { } [] () / \ ' " ` ~ , ; : . < > – A dictionary word (for example, VMware1!)
SDDC Manager super user (vcf)	<ol style="list-style-type: none"> 1. Minimum length 15 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? 3. Must not include: <ul style="list-style-type: none"> – * { } [] () / \ ' " ` ~ , ; : . < > – A dictionary word (for example, VMware1!)

Table continued on next page

Continued from previous page

Password	Requirements
SDDC Manager local account (admin@local)	<ol style="list-style-type: none"> Length 12-127 characters Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? Must not include: * { } [] () / \ ' " ` ~ , ; : . < >

Hosts and Networks Worksheet

The Hosts and Networks worksheet specifies the details for all networks and hosts. This information is configured on the appropriate VMware Cloud Foundation components.

Management Domain Networks

This section covers the VLANs, gateways, MTU, and expected IP ranges and subnet mask for each network you have configured on the Top of Rack switches in your environment.

With VMware Cloud Foundation 5.1 and later, you have the ability to create separate distributed port groups for management VM (for example, vCenter Server and NSX Manager) traffic and ESXi host management traffic.

- If you enter information for the VM Management Network, VMware Cloud Foundation creates a distributed port group for the VM Management Network using the information you provide.
- If you do not enter information for the VM Management Network, VMware Cloud Foundation still creates a distributed port group for VM Management Network, but uses the Management Network information (gateway, VLAN, MTU).

Network Type	VLAN	Portgroup Name	CIDR Notation	Gateway	MTU
VM Management Network	Enter the VLAN ID. The VLAN ID can be between 0 and 4094.	Enter a portgroup name.	Enter the CIDR notation for the network.	Enter the gateway IP for network.	Enter MTU for the network. The MTU can be between 1500 and 9000.

Table continued on next page

Continued from previous page

Network Type	VLAN	Portgroup Name	CIDR Notation	Gateway	MTU
Management Network	NOTE The VLAN ID for Uplink 1 and Uplink 2 Networks must be unique and not used by any other network type.				
vMotion Network					
vSAN Network					

Virtual Networking

The deployment parameter workbook provides three vSphere Distributed Switch profiles that allow you to perform bring-up of hosts with two or four pNICs and to create up to two vSphere Distributed Switches for isolating VMkernel traffic. The information that you are required to provide depends on the profile that you select.

NOTE

You can use the VMware Cloud Foundation API to perform bring-up with other combinations of vSphere Distributed Switches and pNICs that are not available using the vSphere Distributed Switch profiles.

vSphere Distributed Switch Profile	Description
Profile 1	<ul style="list-style-type: none"> • One vSphere Distributed Switch (vDS): Traffic for Management, vMotion, vSAN, and Host Overlay networks using specified pNICs. • Two or four physical NICs (pNICs)
Profile 2	<ul style="list-style-type: none"> • Two vSphere Distributed Switches (vDS) • Four physical NICs (pNICs) • Primary vDS: Traffic for Management, vMotion, and Host Overlay networks using specified pNICs. • Secondary vDS: Traffic for vSAN network using specified pNICs.
Profile 3	<ul style="list-style-type: none"> • Two vSphere Distributed Switches (vDS) • Four physical NICs (pNICs) • Primary vDS: Traffic for Management, vMotion, and vSAN networks using specified pNICs. • Secondary vDS: Traffic for Host Overlay network using specified pNICs.

After you select a vSphere Distributed Switch Profile, enter the required information for that profile.

Primary vSphere Distributed Switch - Name	Enter a name for the primary vSphere Distributed Switch (vDS). You can modify the portgroup names of the management domain networks to make it clear which vDS each network uses.
Primary vSphere Distributed Switch - pNICs	Select the physical NICs to assign to the primary vDS.
Primary vSphere Distributed Switch - MTU Size	Enter the MTU size for the primary vDS. Default value is 9000.
Primary vSphere Distributed Switch - Transport Zone Type	Select Overlay or VLAN.
Secondary vSphere Distributed Switch - Name	Enter a name for the secondary vSphere Distributed Switch (vDS). You can modify the portgroup names of the management domain networks to make it clear which vDS each network uses. NOTE If you are not creating a secondary vDS, enter n/a.
Secondary vSphere Distributed Switch - Transport Zone Type	Select Overlay or VLAN.
Secondary vSphere Distributed Switch - pNICs	Select the physical NICs to assign to the secondary vDS.
Secondary vSphere Distributed Switch - MTU Size	Enter the MTU size for the secondary vDS. Default value is 9000.

Management Domain ESXi Hosts

Specify the IP addresses of the ESXi hosts for the management domain. In a standard deployment, only four hosts are required in the management domain. VMware Cloud Foundation can also be deployed with a consolidated architecture. In a consolidated deployment, all workloads are deployed in the management domain instead of to separate workload domains. As such, additional hosts may be required to provide the capacity needed. In this section, only enter values for the number of hosts desired in the management domain.

Host Name	IP Address
Enter host names for each of the four ESXi hosts.	Enter IP Address for each of the four ESXi hosts.

Inclusion Ranges

Specify IP inclusion ranges for the vSAN and vMotion networks of the management domain. IP addresses from the specified range are automatically assigned to hosts. Ensure that the IP ranges include sufficient IP addresses for the initial deployment. The number of IP addresses must be at least equal to the number of hosts deployed as part of VMware Cloud Foundation.

As an example, if you specify the range start value as 192.168.1.1 and end as 192.168.1.20, a total of 20 IP addresses would be used.

Do not use special IP addresses, such as the network or broadcast address.

IPs for the vMotion range must be part of the VLAN configured with the vMotion portgroup. IPs for the vSAN range must be part of the VLAN configured for the vSAN portgroup. All IPs within the range must be available for use or IP conflicts will occur. It is a good practice to validate this prior to starting a deployment.

Table 166: Input Required

Network	Start IP	End IP
vMotion	Enter start of IP address range for vMotion network.	Enter end of IP address range.
VSAN	Enter start of IP address range for vMotion network.	Enter end of IP address range.

ESXi Host Security Thumbprints

If you want bring-up to validate the SSH fingerprint and SSL thumbprints of the ESXi hosts before connecting to them to reduce the chance of Man In The Middle (MiTM) attack, select **Yes** in the **Validate Thumbprints** field.

If you set **Validate Thumbprints** to **Yes**, follow the steps below.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the navigation pane, click **Manage** and click the **Services** tab.
3. Select the **TSM-SSH** service and click **Start** if not started.
4. Connect to the VMware Cloud Builder appliance using an SSH client such as Putty.
5. Enter the admin credentials you provided when you deployed the VMware Cloud Builder appliance.
6. Retrieve the SSH fingerprint by entering the following command replacing *hostname* with the FQDN of your host:

```
ssh-keygen -lf <(ssh-keyscan hostname 2>/dev/null
```
7. Retrieve the SSL thumbprint by entering the following command replacing *hostname* with the FQDN of your host:

```
openssl s_client -connect hostname:443 < /dev/null 2> /dev/null | openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```
8. In the VMware Host Client, select the **TSM-SSH** service for the ESXi host and click **Stop**.
9. Repeat for each ESXi host and then enter the information in the deployment parameter workbook.

NSX Host Overlay Network

By default, VMware Cloud Foundation uses DHCP for the management domain Host Overlay Network TEPs. For this option, a DHCP server must be configured on the NSX host overlay (Host TEP) VLAN of the management domain. When NSX creates TEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.

For the management domain and VI workload domains with uniform L2 clusters, you can choose to use static IP addresses instead. Make sure the IP range includes enough IP addresses for the number of hosts that will use the static IP Pool. The number of IP addresses required depends on the number of pNICs on the ESXi hosts that are used for the vSphere Distributed Switch that handles host overlay networking. For example, a host with four pNICs that uses two pNICs for host overlay traffic requires two IP addresses in the static IP pool..

Table 167: DHCP Settings

Parameter	Value
VLAN ID	Enter a VLAN ID for the NSX host overlay network. The VLAN ID can be between 0 and 4094.
Configure NSX Host Overlay Using a Static IP Pool	Select No to use DHCP.

Table 168: Static IP Pool Settings

Parameter	Value
VLAN ID	Enter a VLAN ID for the NSX host overlay network. The VLAN ID can be between 0 and 4094.
Configure NSX Host Overlay Using a Static IP Pool	Select Yes to use a static IP pool.
Pool Description	Enter a description for the static IP pool.
Pool Name	Enter a name for the static IP pool.
CIDR Notation	Enter CIDR notation for the NSX Host Overlay network.
Gateway	Enter the gateway IP address for the NSX Host Overlay network.
NSX Host Overlay Start IP	Enter the first IP address to include in the static IP pool.
NSX Host Overlay End IP	Enter the last IP address to include in the static IP pool.

Deploy Parameters Worksheet: Existing Infrastructure Details

Your existing DNS infrastructure is used to provide forward and reverse name resolution for all hosts and VMs in the VMware Cloud Foundation SDDC. External NTP sources are also utilized to synchronize the time between the software components.

Table 169: Infrastructure

Parameter	Value
DNS Server #1	Enter IP address of first DNS server.
DNS Server #2	Enter IP address of second DNS server. NOTE If you have only one DNS server, enter n/a in this cell.
NTP Server #1	Enter IP address or FQDN of first NTP server.
NTP Server #2	Enter IP address or FQDN of second NTP server. NOTE If you have only one NTP server, enter n/a in this cell.

Table 170: DNS Zone

Parameter	Value
DNS Zone Name	Enter root domain name for your SDDC management components. NOTE VMware Cloud Foundation expects all components to be part of the same DNS zone.

Table 171: Customer Experience Improvement Program

Parameter	Value
Enable Customer Experience	Select an option to activate or deactivate CEIP across vSphere, NSX, and vSAN during bring-up.

Table continued on next page

Continued from previous page

Parameter	Value
Improvement Program (“CEIP”)	

Table 172: Enable FIPS Security Mode on SDDC Manager

Parameter	Value
Enable FIPS Security Mode on SDDC Manager	<p>Select an option to activate or deactivate FIPS security mode during bring-up. VMware Cloud Foundation supports Federal Information Processing Standard (FIPS) 140-2. FIPS 140-2 is a U.S. and Canadian government standard that specifies security requirements for cryptographic modules. When you enable FIPS compliance, VMware Cloud Foundation enables FIPS cipher suites and components are deployed with FIPS enabled.</p> <p>To learn more about support for FIPS 140-2 in VMware products, see https://www.vmware.com/solutions/security/certifications/fips.</p> <p>NOTE This option is only available for new VMware Cloud Foundation installations and the setting you apply during bring-up will be used for future upgrades. You cannot change the FIPS security mode setting after bring-up.</p>

Deployment Parameters Worksheet: License Keys

Provide licensing information for VMware Cloud Foundation.

1. Select **Yes** or **No** for **License Now**.
2. If you select **Yes**, in the License Keys section, update the red fields with your license keys. Ensure the license key matches the product listed in each row and that the license key is valid for the version of the product listed in the VMware Cloud Foundation BOM. The license key audit during bring-up validates both the format and validity of the key.

NOTE

When using the per-TiB license for vSAN, be aware that VI workload domain components like vCenter and NSX Manager will also consume the TiB capacity.

3. If you select **No**, the VMware Cloud Foundation components are deployed in evaluation mode.

IMPORTANT

After bring-up, you must switch to licensed mode by adding component license keys in the SDDC Manager UI or adding and assigning a solution license key in the vSphere Client. See the *VMware Cloud Foundation Administration Guide* for information about adding component license keys in the SDDC Manager UI. See [Managing vSphere Licenses](#) for more information about adding and applying a solution license key for VMware ESXi and vCenter Server in the vSphere Client. If you are using a solution license key, you must also add a separate VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

Deploy Parameters Worksheet: vSphere Infrastructure

The vSphere infrastructure section of the Deploy Parameters Worksheet details how you want to configure the vCenter Server and its related objects.

This section of the deployment parameter workbook contains sample configuration information, but you can update them with names that meet your naming standards.

NOTE

All host names entries within the deployment parameter workbook expect the short name. VMware Cloud Builder takes the host name and the DNS zone provided to calculate the FQDN value and performs validation prior to starting the deployment. The specified host names and IP addresses must be resolvable using the DNS servers provided, both forward (hostname to IP) and reverse (IP to hostname), otherwise the bring-up process will fail.

Table 173: vCenter Server

Parameter	Host Name	IP Address
vCenter Server	Enter a host name for the vCenter Server.	Enter the IP address for the vCenter Server that is part of the management VLAN. NOTE This is the same VLAN and IP address space where the ESXi management VMKernels reside.
vCenter Server Appliance Size (Default Small)	This parameter defines the size of the vCenter Server to be deployed. Default size is Small. Additional options are: Tiny, Medium, Large, and X-large. See Hardware Requirements for the vCenter Server Appliance.	
vCenter Server Appliance Storage Size	The amount of storage depends on the vCenter Server appliance size. See Storage Requirements for the vCenter Server Appliance.	

Table 174: vCenter Datacenter and Cluster

Parameter	Value
Datacenter Name	Enter a name for the management datacenter.
Cluster Name	Enter a name for the management cluster.
Enable vLCM Cluster Image	Select Yes to use vSphere Lifecycle Manager images for managing the lifecycle of ESXi hosts in the primary cluster of management domain. VMware Cloud Builder extracts a vSphere Lifecycle Manager image from the first ESXi host and applies that image to all the hosts in the cluster. The vSphere Lifecycle Manager image is also imported into SDDC Manager (available at Lifecycle Management > Image Management). NOTE vSAN Express Storage Architecture (ESA) requires vSphere Lifecycle Manager images. Select No to use vSphere Lifecycle Manager baselines for managing the lifecycle of ESXi hosts in the primary cluster of management domain.
Cluster EVC Setting	To enable EVC on the management cluster, select the CPU chipset that should be applied to enhance vMotion compatibility.

Table continued on next page

Continued from previous page

Parameter	Value
	<p>NOTE If you don't want to enable EVC, enter n/a in this cell.</p>

Select the architecture model you plan to use. If you choose **Consolidated**, specify the names for the vSphere resource pools. You do not need to specify resource pool names if you are using the standard architecture model. See *Introducing VMware Cloud Foundation* for more information about these architecture models.

Table 175: vSphere Resource Pools

Parameter	Value
Resource Pool SDDC Management	Specify the vSphere resource pool name for management VMs.
Resource Pool User Edge	Specify the vSphere resource pool name for user deployed NSX VMs in a consolidated architecture.
Resource Pool User VM	Specify the vSphere resource pool name for user deployed workload VMs.

NOTE

Resource pools are created with Normal CPU and memory shares.

Table 176: vSphere Datastore

Parameter	Value
vSAN Datastore Name	Enter vSAN datastore name for your management components.
Enable vSAN Deduplication and Compression	<p>Select Yes to turn on Dedupe and Compression capabilities of vSAN.</p> <p>NOTE This option is only available with vSAN OSA. If you enable vSAN ESA, deduplication and compression settings can be specified in the vSAN storage policies using the vSphere Client.</p>
Enable vSAN-ESA	<p>Select Yes to use vSAN Express Storage Architecture (ESA) for the first cluster in the management domain. After bringup, you can create additional clusters (vSAN ESA or vSAN OSA) in the management domain.</p> <p>NOTE vSAN ESA requires the use of vLCM images and is not supported with vLCM baselines.</p> <p>vSAN ESA is designed for high-performance NVMe based TLC flash devices and high performance networks. Each host that contributes storage contains a single storage pool</p>

Table continued on next page

Continued from previous page

Parameter	Value
	<p>of four or more flash devices. Each flash device provides caching and capacity to the cluster.</p> <p>Select No to use vSAN Original Storage Architecture (OSA) for the first cluster in the management domain. After bringup, you can create additional clusters (vSAN ESA or vSAN OSA) in the management domain, but you can create vSAN ESA clusters only if the management domain is using vLCM images.</p> <p>For an overview of the differences between vSAN OSA and vSAN ESA, see Building a vSAN Cluster in the vSphere documentation..</p>
Path to HCL JSON File	<p>vSAN ESA requires a current version of the vSAN HCL JSON file to ensure that your ESXi hosts are ESA-compatible.</p> <p>If the VMware Cloud Builder appliance is not able to connect to the internet (either directly or through a proxy server), download the latest vSAN HCL JSON file from https://partnerweb.vmware.com/service/vsan/all.json and copy it to the VMware Cloud Builder appliance.</p> <p>Enter to path to the vSAN HCL JSON file on the VMware Cloud Builder appliance. For example: <code>/opt/vmware/bringup/tmp/all.json</code></p>

If the VMware Cloud Builder appliance does not have direct internet access, you can configure a proxy server to download the vSAN HCL JSON. A recent version of the HCL JSON file is required for vSAN ESA.

Table 177: Proxy Server Configuration

Parameter	Value
Proxy Server Configuration	Select Yes to configure a proxy server.
Proxy Server	Enter the proxy server FQDN or IP address.
Proxy Port	Enter the proxy server port.
Proxy Username	
Proxy Password	
Proxy Transfer Protocol	
HTTPs Proxy Certificate (PEM Encoded)	

Deploy Parameters Worksheet: VMware NSX

The NSX section of the Deploy Parameters Worksheet specifies the details you want to use for deploying VMware NSX components.

Table 178: NSX Management Cluster

Parameter	Value
NSX Management Cluster VIP	<p>Enter the host name and IP address for the NSX Manager VIP.</p> <p>The host name can match your naming standards but must be registered in DNS with both forward and reverse resolution matching the specified IP.</p> <p>NOTE This is the same VLAN and IP address space where the vCenter and ESXi management VMKernels reside.</p>
NSX Virtual Appliance Node #1	Enter the host name and IP address for the first node in the NSX Manager cluster.
NSX Virtual Appliance Node #2	Enter the host name and IP address for the second node in the NSX Manager cluster.
NSX Virtual Appliance Node #3	Enter the host name and IP address for the third node in the NSX Manager cluster.
NSX Virtual Appliance Size	Select the size for the NSX Manager virtual appliances. The default is medium.

Deploy Parameters Worksheet: SDDC Manager

The SDDC Manager section of the Deploy Parameters Worksheet specifies the details for deploying SDDC Manager.

Table 179: SDDC Manager

Parameter	Value
SDDC Manager Hostname	Enter a host name for the SDDC Manager VM.
SDDC Manager IP Address	Enter an IP address for the SDDC Manager VM.
Network Pool Name	Enter the network pool name for the management domain network pool.
Cloud Foundation Management Domain Name	Enter a name for the management domain. This name will appear in Inventory > Workload Domains in the SDDC Manager UI.

Deploy the Management Domain Using ESXi Hosts with External Certificates

VMware Cloud Foundation supports vCenter Server's Custom Certificate Authority mode during bring-up using the VMware Cloud Foundation API. Use this mode if you want to use only external certificates that are signed by a third-party or enterprise CA. In this mode, you are responsible for managing the certificates. You cannot refresh and renew external certificates from the SDDC Manager or vSphere Client.

See [Configure ESXi Hosts with Signed Certificates](#).

To use external ESXi certificates, you must create a custom JSON file for bring-up. You cannot use the deployment parameter workbook.

Deploying the management domain with external ESXi certificates enables Custom Certificate Authority mode, so all future hosts that you add to a workload domain (management or VI) must also use external ESXi certificates.

1. Create a JSON file populated with the bring-up information for your environment.

You can see a sample JSON specification in the [VMware Cloud Foundation API Reference Guide](#).

2. Update the `securitySpec` section, choosing `Custom` for the `esxiCertsMode` and entering your signing CA chain for `certChain`.

For example:

```
"securitySpec" : {
  "esxiCertsMode" : "Custom",
  "rootCaCerts" : [ {
    "alias" : "Rainpole-CA",
    "certChain" : [ "-----BEGIN CERTIFICATE-----
MIIDczCCAlugAwIBAgIQI9xwbTkI9J5GhMffcP5CHDANBgkqhkiG9w0BAQsFADBM
MRIwEAYKCZImizPyLGQBGRYCaW8xGDAWBgoJkiaJk/IsZAEZFgghyYWLucG9sZTEc
MBoGA1UEAxMTcmFpbmBvbGUTZGMwMXJwbClDQTAeFw0yMDAzMzAxNDQ2MTNaFw0y
NTAzMzAxNDU2MTNaMEwxEjAQBgoJkiaJk/IsZAEZFgJpbzEYMBYGCgmSJomT8ixk
ARkVCHJhaW5wb2x1MRwwGgYDVQQDExNyYWLucG9sZS1kYzAxcnBsLUNBMIIBIjAN
BgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEazpwkz7aPlQcfevcCelHc9DPswHkd
kjY96Vh3GvYlesavecy/q/BOvvh3KgLMly8r7cy2cNPO3FANKOfqVdVx3ghfEUyL
g61W9BskAlwryzJRMjhOJJVqvB8CWjy+eCp7MejHGdEud6WdEvK8CaBcPngEg0KM
eLRNLGe8OCw8yY4GTrjU+H7PYQZtyD0kxxy5f48ueaDXat4ENRGCauHEfCoMGfaR
bDue1004diHd900bCym5ggBNX0jhRudNULXPTayZl2ksImV0+QkaVeptQImXfCgb
kgnHQJ5CxK26up7fB5eAsmGLAsJLbnHuM7P9xvV09EvWjFCgLX/oBBDYTQIDAQAB
o1EwTzALBgNVHQ8EBAMCAYYwDwYDVR0TAQH/BAUwAwEB/zAdBgNVHQ4EFgQU7oOq
QBK8yg8mHnAfb+u6/GO0ZUcwEAYJKwYBBAGCNxUBBAMCAQAwDQYJKoZIhvcNAQEL
BQADggEBALYxZGj4vWjFDN1atOUsBx2jrmxbExgMAyRpN1Sc2aj+7vzxHxUW5VbX
x9nc/BfkTiCK6c7Y9VYb+mgjB8z0kNv58sT4ar1yI1ln63VOCoyyLcaFB8HyEJpD
wUhZ4RNPoSijZMpm+M5EuSLfWlhEJo7N8sLqHgVvklDfPbK8fIHbPS5KJwJibbPe
w9UuNRdcxN9hFWKBC0SvfgX+1CJxVdvgfi65rSHPuWinJzrXXdH999DfpDESRzwh
0pqE3GtMct1Nqalp2QJFdahbT+kxj7QWHTjUylSENDHjdlN7a8WH8RGxvEy/97YZ
+crXmxvQ/bAgHk9vcRERbRjfyIs7v88=
-----END CERTIFICATE-----" ] } ] }
```

3. Follow the steps outlined in the [VMware Cloud Foundation API Reference Guide](#) to [deploy the management domain](#).

Troubleshooting VMware Cloud Foundation Deployment

During the deployment stage of VMware Cloud Foundation you can use log files and the Supportability and Serviceability (SoS) Tool to help with troubleshooting.

VMware Cloud Builder Log Files

VMware Cloud Builder contains various log files for different components of the system.

VMware Cloud Builder has a number of components which are used during the bring-up process, each component generates a log file which can be used for the purpose of troubleshooting. The components and their purpose are:

- **JsonGenerator:** Used to convert the deployment parameter workbook into the required configuration file (JSON) that is used by the Bringup Validation Service and Bringup Service.
- **Bringup Service:** Used to perform the validation of the configuration file (JSON), the ESXi hosts and infrastructure where VMware Cloud Foundation will be deployed, and to perform the deployment and configuration of the management domain components and the first cluster.
- **Supportability and Serviceability (SoS) Utility:** A command line utility for troubleshooting deployment issues.

The following table describes the log file locations:

Component	Log Name	Location
JsonGenerator	<code>jsongenerator-timestamp</code>	<code>/var/log/vmware/vcf/sddc-support/</code>
Bringup Service	<code>vcf-bringup.log</code>	<code>/var/log/vmware/vcf/bringup/</code>
	<code>vcf-bringup-debug.log</code>	<code>/var/log/vmware/vcf/bringup/</code>
	<code>rest-api-debug.log</code>	<code>/var/log/vmware/vcf/bringup/</code>
SoS Utility	<code>sos.log</code>	<code>/var/log/vmware/vcf/sddc-support/sos-timestamp/</code>

Using the SoS Utility on VMware Cloud Builder

You can run the Supportability and Serviceability (SoS) Utility on the VMware Cloud Builder appliance to generate a support bundle, which you can use to help debug a failed bring-up of VMware Cloud Foundation.

NOTE

After a successful bring-up, you should only run the SoS Utility on the SDDC Manager appliance. See [Supportability and Serviceability \(SoS\) Utility](#) in the *VMware Cloud Foundation Administration Guide*.

The SoS Utility is not a debug tool, but it does provide health check operations that can facilitate debugging a failed deployment.

To run the SoS Utility in VMware Cloud Builder, SSH in to the VMware Cloud Builder appliance using the `admin` administrative account, then enter `su` to switch to the root user, and navigate to the `/opt/vmware/sddc-support` directory and type `./sos` followed by the options required for your desired operation.

```
./sos --option-1--option-2... --option-n
```

SoS Utility Help Options

Use these options to see information about the SoS tool itself.

Option	Description
--help -h	Provides a summary of the available SoS tool options
--version -v	Provides the SoS tool's version number.

SoS Utility Generic Options

These are generic options for the SoS Utility.

Option	Description
--configure-sftp	Configures SFTP for logs.
--debug-mode	Runs the SoS tool in debug mode.
--force	Allows SoS operations from the VMware Cloud Builder appliance after bring-up. NOTE In most cases, you should not use this option. Once bring-up is complete, you can run the SoS Utility directly from the SDDC Manager appliance.
--history	Displays the last twenty SoS operations performed.
--log-dir <i>LOGDIR</i>	Specifies the directory to store the logs.
--log-folder <i>LOGFOLDER</i>	Specifies the name of the log directory.
--setup-json <i>SETUP_JSON</i>	Custom setup-json file for log collection. SoS prepares the inventory automatically based on the environment where it is running. If you want to collect logs for a pre-defined set of components, you can create a <i>setup.json</i> file and pass the file as input to SoS. A sample JSON file is available on the VMware Cloud Builder in the <i>/opt/vmware/sddc-support/</i> directory.
--skip-known-host-check	Skips the specified check for SSL thumbprint for host in the known host.
--zip	Creates a zipped tar file for the output.

SoS Utility Log File Options

Option	Description
--api-logs	Collects output from APIs.
--cloud-builder-logs	Collects Cloud Builder logs.
--esx-logs	Collects logs from the ESXi hosts only. Logs are collected from each ESXi host available in the deployment.
--no-clean-old-logs	Use this option to prevent the tool from removing any output from a previous collection run.

Table continued on next page

Continued from previous page

Option	Description
	By default, before writing the output to the directory, the tool deletes the prior run's output files that might be present. If you want to retain the older output files, specify this option.
--no-health-check	Skips the health check executed as part of log collection.
--nsx-logs	Collects logs from the NSX Manager instances only.
--rvc-logs	Collects logs from the Ruby vSphere Console (RVC) only. RVC is an interface for ESXi and vCenter. NOTE If the Bash shell is not enabled in vCenter, RVC log collection will be skipped . NOTE RVC logs are not collected by default with ./sos log collection.
--sddc-manager-logs	Collects logs from the SDDC Manager only.
--test	Collects test logs by verifying the files.
--vc-logs	Collects logs from the vCenter Server instances only. Logs are collected from each vCenter server available in the deployment.
--vm-screenshots	Collects screen shots from all VMs.

SoS Utility JSON Generator Options

The JSON generator options within the SoS Utility provide a method to execute the creation of the JSON file from a completed deployment parameter workbook. To run the JSON generator, you must provide, as a minimum, a path to the deployment parameter workbook and the design type using the following syntax:

```
./sos --jsongenerator --jsongenerator-input JSONGENERATORINPUT --jsongenerator-design JSONGENERATORDESIGN
```

Option	Description
--jsongenerator	Invokes the JSON generator utility.
--jsongenerator-input <i>JSONGENERATORINPUT</i>	Specify the path to the input file to be used by the JSON generator utility. For example: /tmp/vcf-ems-deployment-parameter.xlsx.
--jsongenerator-design <i>JSONGENERATORDESIGN</i>	Use vcf-ems for VMware Cloud Foundation.
--jsongenerator-supress	Supress confirmation to force cleanup directory. (optional)
--jsongenerator-logs <i>JSONGENERATORLOGS</i>	Set the directory to be used for logs. (optional)

SoS Utility Health Check Options

The SoS Utility can be used to perform health checks on various components or services, including connectivity, compute, and storage.

NOTE

The health check options are primarily designed to run on the SDDC Manager appliance. Running them on the VMware Cloud Builder appliance requires the `--force` parameter, which instructs the SoS Utility to identify the SDDC Manager appliance deployed by VMware Cloud Builder during the bring-up process, and then execute the health check remotely. For example:

```
./sos --health-check --force
```

Option	Description
<code>--certificate-health</code>	Verifies that the component certificates are valid (within the expiry date).
<code>--connectivity-health</code>	Performs a connectivity health check to inspect whether the different components of the system such as the ESXi hosts, vCenter Servers, NSX Manager VMs, and SDDC Manager VM can be pinged.
<code>--compute-health</code>	Performs a compute health check.
<code>--general-health</code>	Verifies ESXi entries across all sources, checks the Postgres DB operational status for hosts, checks ESXi for error dumps, and gets NSX Manager and cluster status.
<code>--get-host-ips</code>	Returns server information.
<code>--health-check</code>	Performs all available health checks.
<code>--ntp-health</code>	Verifies whether the time on the components is synchronized with the NTP server in the VMware Cloud Builder appliance.
<code>--services-health</code>	Performs a services health check to confirm whether services are running
<code>--run-vsan-checks</code>	Runs proactive vSAN tests to verify the ability to create VMs within the vSAN disks.

Sample Output

The following text is a sample output from an `--ntp-health` operation.

```
root@cloud-builder [ /opt/vmware/sddc-support ]# ./sos --ntp-health --skip-known-host --force
```

```
Welcome to Supportability and Serviceability(SoS) utility!
```

```
User passed --force flag, Running SOS from Cloud Builder VM, although Bringup is completed and SDDC Manager is available. Please expect failures with SoS operations.
```

```
Health Check : /var/log/vmware/vcf/sddc-support/healthcheck-2020-02-11-23-03-53-24681
```

```
Health Check log : /var/log/vmware/vcf/sddc-support/healthcheck-2020-02-11-23-03-53-24681/sos.log
```

```
SDDC Manager : sddc-manager.vrack.vsphere.local
```

```
NTP : GREEN
```

```
+-----+-----+-----+-----+
| SL# |           Area           | Title | State |
+-----+-----+-----+-----+
|  1  | ESXi : esxi-1.vrack.vsphere.local | ESX Time | GREEN |
```

```

| 2 | ESXi : esxi-2.vrack.vsphere.local | ESX Time | GREEN |
| 3 | ESXi : esxi-3.vrack.vsphere.local | ESX Time | GREEN |
| 4 | ESXi : esxi-4.vrack.vsphere.local | ESX Time | GREEN |
| 5 | vCenter : vcenter-1.vrack.vsphere.local | NTP Status | GREEN |
+-----+-----+-----+-----+

```

Legend:

GREEN - No attention required, health status is NORMAL

YELLOW - May require attention, health status is WARNING

RED - Requires immediate attention, health status is CRITICAL

Health Check completed successfully for : [NTP-CHECK]

The following text is sample output from a `--vm-screenshots` log collection operation.

```

root@cloud-builder [ /opt/vmware/sddc-support ]# ./sos --vm-screenshots
--skip-known-host --force

```

Welcome to Supportability and Serviceability(SoS) utility!

User passed `--force` flag, Running SOS from Cloud Builder VM, although Bringup is completed and SDDC Manager is available. Please expect failures with SoS operations.

Logs : /var/log/vmware/vcf/sddc-support/sos-2018-08-24-10-50-20-8013

Log file : /var/log/vmware/vcf/sddc-support/sos-2018-08-24-10-50-20-8013/sos.log

Log Collection completed successfully for : [VMS_SCREENSHOT]

VMware Cloud Foundation Glossary

In VMware Cloud Foundation, you perform specific operations and use unique constructs for automated SDDC deployment and maintenance.

Term	Description
availability zone	A collection of infrastructure components. Each availability zone is isolated from the other availability zones to prevent the propagation of failure or outage across the data center.

Table continued on next page

Continued from previous page

Term	Description
	In VMware Cloud Foundation, you implement availability of workloads across availability zones by using vSAN stretched clusters.
Application virtual networks (AVNs)	Virtual networks backed by overlay or VLAN NSX segments using the encapsulation protocol of VMware NSX. An AVN uses a single IP address space to span across data centers.
bring-up	Deployment and initial configuration of a VMware Cloud Foundation system. During the bring-up process, the management domain is created and the VMware Cloud Foundation software stack is deployed on the management domain.
commission a host	Adding a host to VMware Cloud Foundation inventory. The host becomes unassigned.
dirty host	A host that has been removed from a cluster in a workload domain. A dirty host cannot be assigned to another workload domain until it is decommissioned, re-imaged, and commissioned again.
decommission a host	Removing an unassigned host from the VMware Cloud Foundation inventory. SDDC Manager does not manage decommissioned hosts.
NSX Edge cluster	A logical grouping of NSX Edge nodes. These nodes run on a vSphere cluster, and provide north-south and east-west routing and network services for the management or VI workload domain.
free pool	Hosts in the VMware Cloud Foundation inventory that are not assigned to a workload domain.
host	A server that is imaged with the ESXi software.
install bundle	Contains software for VI workload domains and VMware Aria Suite Lifecycle. You can use an install bundle to deploy later versions of the software components in a new VI workload domain than the versions in the Bill of Materials for VMware Cloud Foundation.
inventory	Logical and physical entities managed by VMware Cloud Foundation.
Kubernetes - Workload Management	With Kubernetes - Workload Management, you can deploy and operate the compute, networking, and storage infrastructure for vSphere IaaS Control Plane workloads. A vSphere IaaS Control Plane workload is an application with containers running inside vSphere pods, regular VMs, or Tanzu Kubernetes clusters.
Lifecycle Manager (LCM)	Automates patching and upgrading of the software stack.
management domain	One or more vSphere clusters of physical hosts that contain the management component VMs, such as vCenter Server, NSX Manager cluster, management NSX Edge cluster, SDDC Manager, and so on. The management domain supports only vSAN storage.
network pool	Automatically assigns static IP addresses to vSAN and vMotion VMkernel ports so that you don't need to enter IP addresses manually when creating a VI workload domain or adding a host or cluster to a workload domain.
update bundle	Contains software to update the VMware Cloud Foundation components in your management or VI workload domain.
principal storage	Required for each vSphere cluster, containing the data of the virtual machines in the cluster. For the management domain, only vSAN principal storage is supported. For a VI workload domain, you set the principal storage when creating the domain or when adding a cluster to the domain. You cannot change the principal storage later. See also <i>supplemental storage</i> .
SDDC Manager	A software component that provisions, manages, and monitors the logical and physical resources of a VMware Cloud Foundation system. SDDC Manager provides the user

Table continued on next page

Continued from previous page

Term	Description
	interface for managing VMware Cloud Foundation, CLI-based administrator tools, and an API for further automation.
server	A bare-metal server in a physical rack. After imaging, it is referred to as a host.
supplemental storage	Extends the capacity of the workload domain for hosting more virtual machines or storing supporting data, such as backups. You can add or remove supplemental storage to clusters in the management or VI workload domain at any time.
unassigned host	A host in the free pool that does not belong to a workload domain.
vSphere Lifecycle Manager (vLCM)	A vCenter Server service, which is integrated with VMware Cloud Foundation, that enables centralized and simplified life cycle management of ESXi hosts.
virtual infrastructure (VI) workload domain	One or more vSphere clusters that contain customer workloads. VMware Cloud Foundation scales and manages the life cycle of each VI workload domain independently. The vCenter Server instance and NSX Manager cluster for a VI workload domain are physically located in the management domain, while the NSX edge nodes - on the VI workload domain.
vSphere Lifecycle Manager baseline	A grouping of multiple bulletins. You can attach a baseline to an ESXi host and check the compliance of the host against the associated baseline. According to the type of content, baselines are patch baselines, extension baselines, and upgrade baselines. SDDC Manager creates the required baseline and baseline group for updating a cluster in a workload domain.
vSphere Lifecycle Manager image	A precise description of the software, components, vendor add-ons, and firmware to run on an ESXi host. You set up a single image and apply it to all hosts in a cluster, thus ensuring cluster-wide host image homogeneity.
workload domain	<p>A policy-based resource container with specific availability and performance attributes that combines vSphere, storage (vSAN, NFS, VMFS on FC, or vVols) and networking (VMware NSX) into a single consumable entity. A workload domain can be created, expanded, and deleted as part of the SDDC life cycle operations. It can contain clusters of physical hosts with a corresponding vCenter Server instance to manage them.</p> <p>VMware Cloud Foundation supports two types of workload domains - the management domain and one or more VI workload domains.</p>

Administration Guide

Understand how to administer and manage your VMware Cloud Foundation system, including password and certificate management, provisioning and configuring workload domains, and deploying additional components, such as vRealize Suite Lifecycle Manager.

About the VMware Cloud Foundation Administration Guide

The *VMware Cloud Foundation Administration Guide* provides information about managing a VMware Cloud Foundation™ system, including managing the system's virtual infrastructure, managing users, configuring, upgrading, and monitoring the system.

Intended Audience

The *VMware Cloud Foundation Administration Guide* is intended for cloud architects, infrastructure administrators, and cloud administrators who are familiar with and want to use VMware software to quickly import, deploy, and manage a software-defined data center (SDDC). The information in this document is written for experienced data center system administrators who are familiar with:

- Concepts of virtualization, software-defined data centers, and virtual infrastructure (VI)
- VMware virtualization technologies, such as VMware ESXi™, the hypervisor
- Software-defined networking using VMware NSX®
- Software-defined storage using VMware vSAN™
- Networking concepts such as Layer-2, Layer-3, and Border Gateway Protocol (BGP).

Related Publications

The *Getting Started with VMware Cloud Foundation* document provides a high-level overview of the VMware Cloud Foundation product.

The *Planning and Preparation Workbook* provides detailed information about the software, tools, and external services that are required for VMware Cloud Foundation.

The *VMware Cloud Foundation Deployment Guide* provides information about installing ESXi software on VMware Cloud Foundation servers and deploying the management domain using the VMware Cloud Builder appliance.

The *VMware Cloud Foundation Lifecycle Management* document describes how to manage the life cycle of a VMware Cloud Foundation environment.

Administering VMware Cloud Foundation

As an SDDC administrator, you use the information in the *VMware Cloud Foundation Administration* document to understand how to administer and operate your VMware Cloud Foundation system.

An administrator of a VMware Cloud Foundation system performs tasks such as:

- Manage certificates and passwords.
- Add capacity to your system.
- Configure and provision workload domains.
- Manage provisioned workload domains.
- Troubleshoot issues and prevent problems across the physical and virtual infrastructure.
- Perform lifecycle management of the VMware Cloud Foundation software components.

See the *Introducing VMware Cloud Foundation* document for a high-level overview of the VMware Cloud Foundation product and the *VMware Cloud Foundation Deployment Guide* for information on deploying the product.

VMware Software Components Deployed by VMware Cloud Foundation

VMware Cloud Foundation is designed and built to deploy specific VMware products using the SDDC Manager appliance.

For the exact version numbers of the VMware products that are supported, refer to the *Release Notes* document for your VMware Cloud Foundation version. If the system has been updated after the initial bring-up process using the lifecycle management features, see "View Upgrade History" in the *VMware Cloud Foundation Lifecycle Management Guide* for details on how to view the versions of the VMware software components that are running in your VMware Cloud Foundation instance.


You can find product-specific documentation for the following VMware software products and components at the [Broadcom Tech Docs portal](#):

- vSphere (vCenter Server and ESXi)
- VMware vSAN
- VMware NSX
- VMware Aria Suite

Web Interfaces Used to Administer VMware Cloud Foundation

SDDC Manager provides a web-based user interface where you can manage your VMware Cloud Foundation instance. This user interface provides centralized access to and an integrated view of the physical and virtual infrastructure of your system.

In addition to using the SDDC Manager UI, you can use the following user interfaces for administration tasks involving their associated VMware software components that are part of a VMware Cloud Foundation instance. All of these interfaces run in a web browser, and you can launch them from within the SDDC Manager UI.

Launch links are typically identified in the SDDC Manager UI by the launch icon: .

VMware SDDC Web Interfaces	Launch Link Location in the SDDC Manager UI
vSphere Client	<ol style="list-style-type: none"> 1. In the navigation pane, click Inventory > Workload Domains. 2. Click View Details for a workload domain. 3. In the Domain column, click the domain name. 4. Click the Services tab. 5. Click the vCenter Server launch link.
NSX Manager UI	<ol style="list-style-type: none"> 1. In the navigation pane, click Inventory > Workload Domains. 2. Click View Details for a workload domain. 3. In the Domain column, click the domain name. 4. Click the Services tab. 5. Click the NSX Cluster launch link.
VMware Host Client	<ol style="list-style-type: none"> 1. In the navigation pane, click Inventory > Hosts.. 2. In the FQDN column, click the host FQDN. 3. Click Actions > Open in VMware Host Client.

Getting Started with SDDC Manager

You use SDDC Manager to perform administration tasks on your VMware Cloud Foundation instance. The SDDC Manager UI provides an integrated view of the physical and virtual infrastructure and centralized access to manage the physical and logical resources.

You work with the SDDC Manager UI by loading it in a web browser. For the list of supported browsers and versions, see the *Release Notes*.

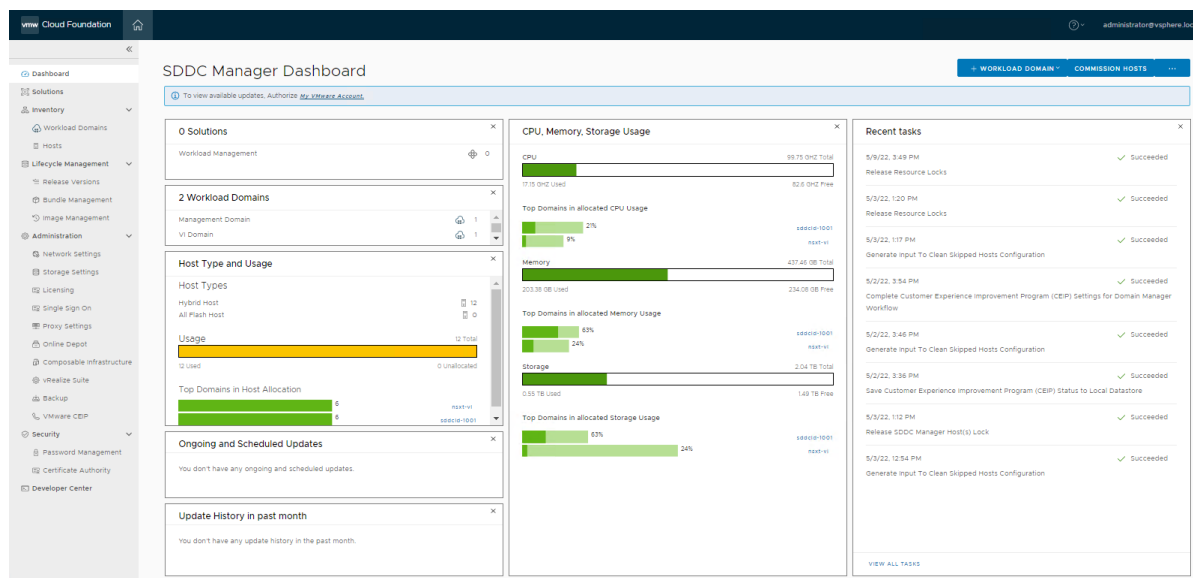
Log in to the SDDC Manager User Interface

Connect to the SDDC Manager appliance by logging into the SDDC Manager UI using a supported web browser.

To log in, you need the SDDC Manager IP address or FQDN and the password for the single-sign on user (for example `administrator@vsphere.local`). You added this information to the deployment parameter workbook before bring-up.

- In a web browser, type one of the following.
 - `https://FQDN` where *FQDN* is the fully-qualified domain name of the SDDC Manager appliance.
 - `https://IP_address` where *IP_address* is the IP address of the SDDC Manager appliance.
- Log in to the SDDC Manager UI with vCenter Server Single Sign-On user credentials.

You are logged in to SDDC Manager UI and the Dashboard page appears in the web browser.



Guided SDDC Manager Onboarding

VMware Cloud Foundation includes an onboarding dashboard to help you with configuring a healthy SDDC Manager environment.

This dashboard appears when you log into SDDC Manager. It provides a walk-through for initial configuration, including the recommended order for completing each task. After completing the walk-through, a banner at the top of the screen offers a tour of the SDDC Manager UI.

You can skip sections and exit out of the guided setup at any point. This dashboard automatically shows unless you click "Don't show onboarding screen again" and close the page. Clicking this option also prevents the optional guided tour from automatically displaying in the future.

Use the Help icon in the upper-right corner of the page to later access the onboarding dashboard and guided tour.

Tour of the SDDC Manager User Interface

The SDDC Manager UI provides a single point of control for managing and monitoring your VMware Cloud Foundation instance and for provisioning workload domains.

You use the navigation bar to move between the main areas of the user interface.

Navigation Bar

The navigation bar is available on the left side of the interface and provides a hierarchy for navigating to the corresponding pages.

Category	Functional Areas
Dashboard	<p>The Dashboard provides the high-level administrative view for SDDC Manager in the form of widgets. There are widgets for Solutions; Workload Domains; Host Types and Usage; Ongoing and Scheduled Updates; Update History; CPU, Memory, Storage Usage; and Recent Tasks.</p> <p>You can control the widgets that are displayed and how they are arranged on the dashboard.</p> <ul style="list-style-type: none"> To rearrange widgets, click the heading of the widget and drag it to the desired position. To hide a widget, hover the mouse anywhere over the widget to reveal the X in the upper-right corner, and click the X. To add a widget, click the three dots in the upper right corner of the page and select Add New Widgets. This displays all hidden widgets. Select a widget and click Add.
Solutions	<p>Solutions include the following section:</p> <ul style="list-style-type: none"> Kubernetes - Workload Management allows you to start a Workload Management deployment and view Workload Management cluster details.
Inventory	<p>Inventory includes the following sections:</p> <ul style="list-style-type: none"> Workload Domains takes you to the Workload Domains page, which displays and provides access to all workload domains. <p>This page includes summary information about all workload domains, including domain type, storage usage, configuration status, owner, clusters, hosts and update availability. It also displays CPU, memory, and storage utilization for each workload domain, and collectively across all domains.</p> <ul style="list-style-type: none"> Hosts takes you to the Hosts page, which displays and provides access to current hosts and controls for managing hosts. <p>This page includes detailed information about all hosts, including FQDN, host IP, network pool, configuration status, host state, cluster, and storage type. It also displays CPU and memory utilization for each host, and collectively across all hosts.</p>
Lifecycle Management	<p>Lifecycle Management includes the following sections:</p>

Table continued on next page

Continued from previous page

Category	Functional Areas
	<ul style="list-style-type: none"> • Release Versions displays the versions in your environment and the associated component versions in that release. • Bundle Management displays the available install, update, and upgrade bundles for your environment, and your bundle download history. <p style="text-align: center;">NOTE To access bundles, you must be logged in to your Broadcom Support Portal account through the Administration > Depot Settings page.</p> <ul style="list-style-type: none"> • Image Management allows you to import a vSphere Lifecycle Manager cluster image from vCenter Server and view the available images. This is an alternative way of managing the life cycle of ESXi hosts.
Administration	<p>Administration includes the following sections:</p> <ul style="list-style-type: none"> • Network Settings allows you to configure, view, and manage network pool settings. You can create new network pools, and view and modify existing network pools. You can also use Network Settings to update the DNS and NTP servers that VMware Cloud Foundation uses. • Storage Settings allows you to add new VASA providers, view, edit, and delete existing VASA providers. • Licensing allows you to manage VMware product licenses. You can also add licenses for the component products in your VMware Cloud Foundation deployment. • Single Sign On allows you to manage VMware Cloud Foundation users and groups, including adding users and groups and assigning roles. You can also configure identity providers for VMware Cloud Foundation. • Proxy Settings allows you to configure a proxy server to download install and upgrade bundles from the VMware Depot. • Depot Settings allows you to log in to your Broadcom Support Portal account to download install and upgrade bundles. • VMware Aria Suite allows you to deploy VMware Aria Suite Lifecycle and configure connections between workload domains and VMware Aria Suite products. • Backup allows you to register an external SFTP server with SDDC Manager for backing up SDDC Manager and NSX Managers. You can also configure the backup schedule for SDDC Manager. • VMware CEIP to join or leave the VMware Customer Experience Improvement Program.

Table continued on next page

Continued from previous page

Category	Functional Areas
Security	<ul style="list-style-type: none"> • Password Management allows password management actions, such as rotation, updates and remediation. • Certificate Authority allows you to integrate with your Microsoft Certificate Authority Server.
Developer Center	<p>The VMware Cloud Foundation Developer Center includes the following sections:</p> <ul style="list-style-type: none"> • Overview: API reference documentation. Includes information and steps for all the Public APIs supported by VMware Cloud Foundation. • API Explorer: Lists the APIs and allows you to invoke them directly on your VMware Cloud Foundation system.

Log out of the SDDC Manager User Interface

Log out of the SDDC Manager UI when you have completed your tasks.

1. In the SDDC Manager UI, click the logged-in account name in the upper right corner.
2. Click **Log out**.


Configure the Customer Experience Improvement Program Settings for VMware Cloud Foundation

VMware Cloud Foundation participates in the VMware Customer Experience Improvement Program (CEIP). You can choose to activate or deactivate CEIP for your VMware Cloud Foundation instance.

The Customer Experience Improvement Program provides VMware with information that allows VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects technical information about your organization's use of the VMware products and services regularly in association with your organization's VMware license keys. This information does not personally identify any individual. For additional information regarding the CEIP, refer to the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

You can activate or deactivate CEIP across all the components deployed in VMware Cloud Foundation by the following methods:

- When you log into SDDC Manager for the first time, a pop-up window appears. The **Join the VMware Customer Experience Program** option is selected by default. Deselect this option if you do not want to join CEIP. Click **Apply**.



- You can activate or deactivate CEIP from the Administration tab in the SDDC Manager UI.

1. In the navigation pane, click **Administration** > **VMware CEIP**.
2. To activate CEIP, select the **Join the VMware Customer Experience Improvement Program** option.
3. To deactivate CEIP, deselect the **Join the VMware Customer Experience Improvement Program** option.

Managing Certificates in VMware Cloud Foundation

You can use the SDDC Manager UI to manage certificates in a VMware Cloud Foundation instance, including integrating a certificate authority, generating and submitting certificate signing requests (CSR) to a certificate authority, and downloading and installing certificates.

Starting with VMware Cloud Foundation 5.2.1, you can also manage certificates using the vSphere Client.

This section provides instructions for the SDDC Manager UI to:

- Use OpenSSL as a certificate authority, which is a native option in SDDC Manager.
- Integrate with Microsoft Active Directory Certificate Services.
- Provide signed certificates from another external Certificate Authority.

You can manage the certificates for the following components.

- vCenter Server
- NSX Manager
- VMware Avi Load Balancer (formerly known as NSX Advanced Load Balancer)
- SDDC Manager
- VMware Aria Suite Lifecycle

NOTE

Use VMware Aria Suite Lifecycle to manage certificates for the other VMware Aria Suite components.

NOTE

VMware Cloud Foundation does not manage certificates for ESXi hosts. By default, ESXi hosts use VMCA-signed certificates, but they can also use external CA-signed certificates. If ESXi hosts are using VMCA-signed certificates, VMCA manages the certificates and certificate rotation. If ESXi hosts are using external certificates, you are responsible for managing the certificates. For more information about external certificates, see [Configure ESXi Hosts with Signed Certificates](#).

You replace certificates for the following reasons:

- A certificate has expired or is nearing its expiration date.
- A certificate has been revoked by the issuing certificate authority.
- You do not want to use the default VMCA-signed certificates.
- Optionally, when you create a new workload domain.

It is recommended that you replace all certificates after completing the deployment of the VMware Cloud Foundation management domain. After you create a new VI workload domain, you can replace certificates for the appropriate components as needed.

View Certificate Information

You can view details of an applied certificate for a resource directly through the SDDC Manager UI.

The SDDC Manager UI provides a banner notification for any certificates that are expiring in the next 30 days.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the **Workload Domains** page, from the table, in the domain column click the domain you want to view.
3. On the domain summary page, click the **Certificates** tab.

Summary Services Updates Update History Hosts Clusters Edge Clusters Certificates								
GENERATE CSRS		GENERATE SIGNED CERTIFICATES		INSTALL CERTIFICATES		DOWNLOAD CSR		UPLOAD AND INSTALL CERTIFICATES
<input type="checkbox"/>	Resource Type	Issuer	Issued To	Valid From	Valid Until	Status	Certificate Operation Status	
<input type="checkbox"/>	> vcenter	CA	vcenter-vsan.vrack.vsphere.lo...	Jun 8, 2023	Jun 7, 2025	Active	CSR Generation - NOT STARTED	
<input type="checkbox"/>	> nsx	CA	nsxt-manager-1-cl1.vrack.vsphere.local	Jun 8, 2023	Sep 10, 2025	Active	CSR Generation - NOT STARTED	
<input type="checkbox"/>	> nsx	CA	vip-nsxmanager-cl1.vrack.vsphere.local	Jun 8, 2023	Sep 10, 2025	Active	CSR Generation - NOT STARTED	

This tab lists the certificates for each resource type associated with the workload domain. It displays the following details:

- Resource type
- Issuer, the certificate authority name
- Resource hostname
- Valid From
- Valid Until
- Certificate status: Active, Expiring, or Expired.
- Certificate operation status

4. To view certificate details, expand the resource next to the Resource Type column.

Configure VMware Cloud Foundation to Use Microsoft CA-Signed Certificates

VMware Cloud Foundation supports the ability to manage certificates by integrating with Microsoft Active Directory Certificate Services (Microsoft CA). Before you can perform certificate operations using the SDDC Manager UI you must ensure that the Microsoft Certificate Authority is configured correctly.

Complete the below tasks to manage Microsoft CA-Signed certificates using SDDC Manager.

Prepare Your Microsoft Certificate Authority to Allow SDDC Manager to Manage Certificates

To ensure secure and operational connectivity between the SDDC components, you apply signed certificates provided by a Microsoft Certificate Authority for the SDDC components.

You use SDDC Manager to generate the certificate signing request (CSRs) and request a signed certificate from the Microsoft Certificate Authority. SDDC Manager is then used to install the signed certificates to SDDC components it manages. In order to achieve this the Microsoft Certificate Authority must be configured to allow integration with SDDC Manager.

Install Microsoft Certificate Authority Roles

Install the Certificate Authority and Certificate Authority Web Enrollment roles on the Microsoft Certificate Authority server to facilitate certificate generation from SDDC Manager.

NOTE

When connecting SDDC Manager to Microsoft Active Directory Certificate Services, ensure that Web Enrollment role is installed on the same machine where the Certificate Authority role is installed. SDDC Manager can't request and sign certificates automatically if the two roles (Certificate Authority and Web Enrollment roles) are installed on different machines.

1. Log in to the Microsoft Certificate Authority server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Add roles to Microsoft Certificate Authority server.
 - a) Click **Start > Run**, enter `ServerManager`, and click **OK**.
 - b) From the **Dashboard**, click **Add roles and features** to start the **Add Roles and Features** wizard.
 - c) On the **Before you begin** page, click **Next**.
 - d) On the **Select installation type** page, click **Next**.
 - e) On the **Select destination server** page, click **Next**.
 - f) On the **Select server roles** page, under **Active Directory Certificate Services**, select **Certification Authority** and **Certification Authority Web Enrollment** and click **Next**.
 - g) On the **Select features** page, click **Next**.
 - h) On the **Confirm installation selections** page, click **Install**.

Configure the Microsoft Certificate Authority for Basic Authentication

Configure the Microsoft Certificate Authority with basic authentication to allow SDDC Manager the ability to manage signed certificates.

The Microsoft Certificate Authority and IIS must be installed on the same server.

1. Log in to the Active Directory server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Add Basic Authentication to the Web Server (IIS).
 - a) Click **Start > Run**, enter `ServerManager`, and click **OK**.
 - b) From the **Dashboard**, click **Add roles and features** to start the **Add Roles and Features** wizard.
 - c) On the **Before you begin** page, click **Next**.
 - d) On the **Select installation type** page, click **Next**.
 - e) On the **Select destination server** page, click **Next**.
 - f) On the **Select server roles** page, under **Web Server (IIS) > Web Server > Security**, select **Basic Authentication** and click **Next**.

- g) On the **Select features** page, click **Next**.
- h) On the **Confirm installation selections** page, click **Install**.
3. Configure the certificate service template and CertSrv web site, for basic authentication.
 - a) Click **Start > Run**, enter `Inetmgr.exe` and click **OK** to open the **Internet Information Services Application Server Manager**.
 - b) Navigate to *your_server* > **Sites** > **Default Web Site** > **CertSrv**.
 - c) Under **IIS**, double-click **Authentication**.
 - d) On the **Authentication** page, right-click **Basic Authentication** and click **Enable**.
 - e) In the navigation pane, select **Default Web Site**.
 - f) In the **Actions** pane, under **Manage Website**, click **Restart** for the changes to take effect.

Create and Add a Microsoft Certificate Authority Template

You must set up a certificate template in the Microsoft Certificate Authority. The template contains the certificate authority attributes for signing certificates for the VMware Cloud Foundation components. After you create the template, you add it to the certificate templates of the Microsoft Certificate Authority.

1. Log in to the Active Directory server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Click **Start > Run**, enter `certtmpl.msc`, and click **OK**.
3. In the **Certificate Template Console** window, under **Template Display Name**, right-click **Web Server** and select **Duplicate Template**.
4. In the **Properties of New Template** dialog box, click the **Compatibility** tab and configure the following values.

Setting	Value
Certification Authority	Windows Server 2008 R2
Certificate recipient	Windows 7 / Server 2008 R2

5. In the **Properties of New Template** dialog box, click the **General** tab and enter a name for example, `VMware` in the **Template display name** text box.
6. In the **Properties of New Template** dialog box, click the **Extensions** tab and configure the following.
 - a) Click **Application Policies** and click **Edit**.
 - b) Click **Server Authentication**, click **Remove**, and click **OK**.
 - c) Click **Basic Constraints** and click **Edit**.
 - d) Click the **Enable this extension** check box and click **OK**.
 - e) Click **Key Usage** and click **Edit**.
 - f) Click the **Signature is proof of origin (nonrepudiation)** check box, leave the defaults for all other options and click **OK**.
7. In the **Properties of New Template** dialog box, click the **Subject Name** tab, ensure that the **Supply in the request** option is selected, and click **OK** to save the template.
8. Add the new template to the certificate templates of the Microsoft CA.
 - a) Click **Start > Run**, enter `certsrv.msc`, and click **OK**

- b) In the **Certification Authority** window, expand the left pane, right-click **Certificate Templates**, and select **New > Certificate Template to Issue**.
- c) In the **Enable Certificate Templates** dialog box, select **VMware**, and click **OK**.

Assign Certificate Management Privileges to the SDDC Manager Service Account

Before you can use the Microsoft Certificate Authority and the pre-configured template, it is recommended to configure least privilege access to the Microsoft Active Directory Certificate Services using an Active Directory user account as a restricted service account.

- Create a user account in Active Directory with Domain Users membership. For example, `svc-vcf-ca`.

1. Log in to the Microsoft Certificate Authority server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Configure least privilege access for a user account on the Microsoft Certificate Authority.

- a) Click **Start > Run**, enter `certsrv.msc`, and click **OK**.
- b) Right-click the certificate authority server and click **Properties**.
- c) Click the **Security** tab, and click **Add**.
- d) Enter the name of the user account and click **OK**.
- e) In the **Permissions for** section configure the permissions and click **OK**.

Setting	Value (Allow)
Read	Deselected
Issue and Manage Certificates	Selected
Manage CA	Deselected
Request Certificates	Selected

3. Configure least privilege access for the user account on the Microsoft Certificate Authority Template.

- a) Click **Start > Run**, enter `certtmpl.msc`, and click **OK**.
- b) Right-click the VMware template and click **Properties**.
- c) Click the **Security** tab, and click **Add**.
- d) Enter the `svc-vcf-ca` service account and click **OK**.
- e) In the **Permissions for** section configure the permissions and click **OK**.

Setting	Value (Allow)
Full Control	Deselected
Read	Selected
Write	Deselected
Enroll	Selected

Table continued on next page

Continued from previous page

Setting	Value (Allow)
Autoenroll	Deselected

Configure a Microsoft Certificate Authority in SDDC Manager

You configure a connection between SDDC Manager and a Microsoft Certificate Authority by entering your service account credentials.

- Verify connectivity between SDDC Manager and the Microsoft Certificate Authority Server. See [VMware Ports and Protocols](#).
- Verify that the Microsoft Certificate Authority Server has the correct roles installed on the same machine where the Certificate Authority role is installed. See [Install Microsoft Certificate Authority Roles](#).
- Verify the Microsoft Certificate Authority Server has been configured for basic authentication. See [Configure the Microsoft Certificate Authority for Basic Authentication](#).
- Verify a valid certificate template has been configured on the Microsoft Certificate Authority. See [Create and Add a Microsoft Certificate Authority Template](#).
- Verify least privileged user account has been configured on the Microsoft Certificate Authority Server and Template. See [Assign Certificate Management Privileges to the SDDC Manager Service Account](#).
- Verify that time is synchronized between the Microsoft Certificate Authority and the SDDC Manager appliance. Each system can be configured with a different timezone, but it is recommended that they receive their time from the same NTP source.

1. In the navigation pane, click **Security** > **Certificate Authority**.
2. Click **Edit**.

Configure Certificate Authority EDIT

Certificate Authority Type	<u>Microsoft</u> ▾
CA Server URL ⓘ	<u>https://<url>/certsrv</u>
User Name	<u>User Name</u>
Password	<u>Password</u> 👁
Template Name ⓘ	<u>Template Name</u>

SAVE

3. Configure the settings and click **Save**.

Setting	Value
Certificate Authority Type	Microsoft
CA Server URL	Specify the URL for the issuing certificate authority. This address must begin with <code>https://</code> and end with <code>certsrv</code> . For example, <code>https://ca.rainpole.io/certsrv</code> .
User Name	Enter a least privileged service account. For example, <code>svc-vcf-ca</code> .
Password	Enter the password for the least privileged service account.
Template Name	Enter the issuing certificate template name. You must create this template in Microsoft Certificate Authority. For example, VMware.

- In the **CA Server Certificate Details** dialog box, click **Accept**.

Install Microsoft CA-Signed Certificates using SDDC Manager

Replace the self-signed certificates with signed certificates from the Microsoft Certificate Authority by using SDDC Manager.

- In the navigation pane, click **Inventory** > **Workload Domains**.
- On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
- On the domain summary page, click the **Certificates** tab.

Resource Type	Issuer	Issued To	Valid From	Valid Until	Status	Certificate Operation Status
vcenter	CA	vcenter-1.vrack.vsphere.local	Jun 2, 2023	Jun 1, 2025	Active	CSR Generation - NOT STARTED

- Generate CSR files for the target components.
 - From the table, select the check box for the resource type for which you want to generate a CSR.
 - Click **Generate CSRs**.
 - On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.

Table continued on next page

Continued from previous page

Option	Description
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.
 - e) On the **Summary** dialog, click **Generate CSRs**.
5. Generate signed certificates for each component.
 - a) From the table, select the check box for the resource type for which you want to generate a signed certificate for.
 - b) Click **Generate Signed Certificates**.
 - c) In the **Generate Certificates** dialog box, from the **Select Certificate Authority** drop-down menu, select **Microsoft**.
 - d) Click **Generate Certificates**.
 6. Install the generated signed certificates for each component.
 - a) From the table, select the check box for the resource type for which you want to install a signed certificate.
 - b) Click **Install Certificates**.

Configure VMware Cloud Foundation to Use OpenSSL CA-Signed Certificates

VMware Cloud Foundation supports the ability to manage certificates using OpenSSL configured on the SDDC Manager appliance.

Complete the following tasks to be able to manage OpenSSL-signed certificates issued by SDDC Manager.

Configure OpenSSL-signed Certificates in SDDC Manager

To generate OpenSSL-signed certificates for the VMware Cloud Foundation components you must first configure the certificate authority details.

1. In the navigation pane, click **Security** › **Certificate Authority**.
2. Click **Edit**.
3. Configure the settings and click **Save**.

Configure Certificate Authority

Certificate Authority Type	OpenSSL ▾
Common Name	Not Specified
Organizational Unit	Not Specified
Organization	Not Specified
Locality	Not Specified
State	Not Specified
Country	Please Select Country ▾


Setting	Value
Certificate Authority	OpenSSL
Common Name	Specify the FQDN of the SDDC Manager appliance.
Organizational Unit	Use this field to differentiate between the divisions within your organization with which this certificate is associated.
Organization	Specify the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Specify the city or the locality where your company is legally registered.
State	Enter the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Select the country where your company is registered. This value must use the ISO 3166 country code.

Install OpenSSL-signed Certificates using SDDC Manager

Replace the self-signed certificates with OpenSSL-signed certificates generated by SDDC Manager.

1. In the navigation pane, click **Inventory** › **Workload Domains**.

2. On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
3. On the domain summary page, click the **Certificates** tab.



4. Generate CSR files for the target components.
 - a) From the table, select the check box for the resource type for which you want to generate a CSR.
 - b) Click **Generate CSRs**.
The **Generate CSRs** wizard opens.
 - c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.
You can enter multiple values separated by comma (,), semicolon (;), or space (). For NSX, you can enter the subject alternative name for each node along with the Virtual IP (primary) node.

NOTE

Wildcard subject alternate name, such as *.example.com is not recommended.

- e) On the **Summary** dialog, click **Generate CSRs**.
5. Generate signed certificates for each component.
 - a) From the table, select the check box for the resource type for which you want to generate a signed certificate.

- b) Click **Generate Signed Certificates**.
 - c) In the **Generate Certificates** dialog box, from the **Select Certificate Authority** drop-down menu, select **OpenSSL**.
 - d) Click **Generate Certificates**.
6. Install the generated signed certificates for each component.
- a) From the table, select the check box for the resource type for which you want to install a signed certificate.
 - b) Click **Install Certificates**.

Install Third-Party CA-Signed Certificates Using Server Certificate and Certificate Authority Files

VMware Cloud Foundation supports two ways to install third-party certificates. This procedure describes the new method, which is the default method for VMware Cloud Foundation 4.5.1 and later.

If you prefer to use the legacy method for installing third-party CA-signed certificates, see [Install Third-Party CA-Signed Certificates in VMware Cloud Foundation Using a Certificate Bundle](#).

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
3. On the domain summary page, click the **Certificates** tab.
4. Generate CSR files for the target components.
 - a) From the table, select the check box for the resource type for which you want to generate a CSR.
 - b) Click **Generate CSRs**.
The **Generate CSRs** wizard opens.
 - c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.

You can enter multiple values separated by comma (,), semicolon (;), or space (). For NSX, you can enter the subject alternative name for each node along with the Virtual IP (primary) node.

NOTE

Wildcard subject alternative name, such as *.example.com are not recommended.

- e) On the **Summary** dialog, click **Generate CSRs**.
5. Download and save the CSR files by clicking **Download CSR**.
 6. When the downloads complete, request signed certificates from your third-party Certificate Authority for each .csr.
 7. After you receive the signed certificates, open the SDDC Manager UI and click **Upload and Install**.
 8. In the **Install Signed Certificates** dialog box, select the resource for which you want to install a signed certificate. The drop-down menu includes all resources for which you have generated and downloaded CSRs.
 9. Select a **Source** and enter the required information.

Source	Required Information
Paste Text	<p>Copy and paste the:</p> <ul style="list-style-type: none"> • Server Certificate • Certificate Authority <p>Paste the server certificate and the certificate authority in PEM format (base64-encoded) . For example:</p> <pre>-----BEGIN CERTIFICATE----- <certificate content> -----END CERTIFICATE-----</pre> <p>If the Certificate Authority includes intermediate certificates, it should be in the following format:</p> <pre>-----BEGIN CERTIFICATE----- <Intermediate certificate content> -----END CERTIFICATE----- -----BEGIN CERTIFICATE----- <Root certificate content> -----END CERTIFICATE-----</pre>
File Upload	<p>Click Browse to upload the:</p> <ul style="list-style-type: none"> • Server Certificate • Certificate Authority <p>Files with .crt, .cer, .pem, .p7b and .p7c extensions are supported.</p>
Certificate Chain	<p>Click Browse to upload the certificate chain. Files with .crt, .cer, .pem, .p7b and .p7c extensions are supported.</p>

10. Click **Validate**.
If validation fails, resolve the issues and try again, or click **Remove** to skip the certificate installation.
11. To install a signed certificate for another resource, click **Add Another** and repeat steps 8-10 for each resource.


- Once all signed certificates have been validated successfully, click **Install**.

Install Third-Party CA-Signed Certificates in VMware Cloud Foundation Using a Certificate Bundle

VMware Cloud Foundation supports two ways to install third-party certificates. This procedure describes the legacy method of using a certificate bundle. To use the legacy method, you must modify your preferences and then use this procedure to generate CSRs, sign the CSRs with a third-party CA, and finally upload and install the certificates.

VMware Cloud Foundation 4.5.1 introduces a new method for installing third-party CA-signed certificates. By default, VMware Cloud Foundation use the new method. See [Install Third-Party CA-Signed Certificates Using Server Certificate and Certificate Authority Files](#) for information using the new method. If you prefer to use the legacy method, you must modify your preferences.

- In the SDDC Manager UI, click the logged in user and select **Preferences**.



- Use the toggle to switch to legacy certificate management.

Revert to Legacy Certificate Management



Uploading CA-signed certificates from a third-party Certificate Authority using the legacy method requires that you collect the relevant certificate files in the correct format and then create a single `.tar.gz` file with the contents. It's important that you create the correct directory structure within the `.tar.gz` file as follows:

- The name of the top-level directory must exactly match the name of the workload domain as it appears in the list on the **Inventory > Workload Domains**. For example, `sfo-m01`.
 - The PEM-encoded root CA certificate chain file (must be named `rootca.crt`) must reside inside this top-level directory. The `rootca.crt` chain file contains a root certificate authority and can have `n` number of intermediate certificates.

For example:

```
-----BEGIN CERTIFICATE-----
<Intermediate1 certificate content>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Intermediate2 certificate content>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root certificate content>
```

```
-----END CERTIFICATE-----
```

In the above example, there are two intermediate certificates, *intermediate1* and *intermediate2*, and a root certificate. *intermediate1* must use the certificate issued by *intermediate2* and *intermediate2* must use the certificate issued by Root CA.

- The root CA certificate chain file, intermediate certificates, and root certificate must contain the `Basic Constraints` field with value `CA:TRUE`.
- This directory must contain one sub-directory for each component resource for which you want to replace the certificates.
- Each sub-directory must exactly match the resource hostname of a corresponding component as it appears in the Resource Hostname column in the **Inventory > Workload Domains > Certificates** tab.

For example, `nsxManager.vrack.vsphere.local`, `vcenter-1.vrack.vsphere.local`, and so on.

- Each sub-directory must contain the corresponding `.csr` file, whose name must exactly match the resource as it appears in the Resource Hostname column in the **Inventory > Workload Domains > Certificates** tab.
- Each sub-directory must contain a corresponding `.crt` file, whose name must exactly match the resource as it appears in the Resource Hostname column in the **Inventory > Workload Domains > Certificates** tab. The content of the `.crt` files must end with a newline character.

For example, the `nsxManager.vrack.vsphere.local` sub-directory would contain the `nsxManager.vrack.vsphere.local.crt` file.

- All certificates including `rootca.crt` must be in UNIX file format.
- Additional requirements for NSX certificates:
 - Server certificate (`NSX_FQDN.crt`) must contain the `Basic Constraints` field with value `CA:FALSE`.
 - If the NSX certificate contains HTTP or HTTPS based CRL Distribution Point it must be reachable from the server.
 - The extended key usage (EKU) of the generated certificate must contain the EKU of the CSR generated.

NOTE

All resource and hostname values can be found in the list on the **Inventory > Workload Domains > Certificates** tab.

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
3. On the domain summary page, click the **Certificates** tab.
4. Generate CSR files for the target components.
 - a) From the table, select the check box for the resource type for which you want to generate a CSR.
 - b) Click **Generate CSRs**.
The **Generate CSRs** wizard opens.
 - c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.

Table continued on next page

Continued from previous page

Option	Description
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**. You can enter multiple values separated by comma (,), semicolon (;), or space (). For NSX, you can enter the subject alternative name for each node along with the Virtual IP (primary) node.

NOTE

Wildcard subject alternative name, such as *.example.com are not recommended.

- e) On the **Summary** dialog, click **Generate CSRs**.
- Download and save the CSR files to the directory by clicking **Download CSR**.
 - Complete the following tasks outside of the SDDC Manager UI:
 - Verify that the different .csr files have successfully generated and are allocated in the required directory structure.
 - Request signed certificates from a Third-party Certificate authority for each .csr.
 - Verify that the newly acquired .crt files are correctly named and allocated in the required directory structure.
 - Create a new .tar.gz file of the directory structure ready for upload to SDDC Manager. For example:
 <domain name>.tar.gz.
 - Click **Upload and Install**.
 - In the **Upload and Install Certificates** dialog box, click **Browse** to locate and select the newly created <domain name>.tar.gz file and click **Open**.
 - Click **Upload**.
 - If the upload is successful, click **Install Certificate**. The Certificates tab displays a status of Certificate Installation is in progress.

Add a Trusted Certificate to the SDDC Manager Trust Store

If you replaced the certificate for a VMware Cloud Foundation component outside of SDDC Manager then you must add the new certificate to the SDDC Manager trust store.

This functionality is available in VMware Cloud Foundation 4.5.1 and later.

Replacing the certificate for a VMware Cloud Foundation component outside of SDDC Manager results in an error in the SDDC Manager UI.



You can add the trusted certificate to the SDDC Manager trust store using the VMware Cloud Foundation API or the SDDC Manager UI. This procedure describes using the SDDC Manager UI. Using the SDDC Manager UI adds the certificate to the trust store for outbound communications.

1. Click **review** in the error message in the SDDC Manager UI.
In the SDDC Manager UI, click **Inventory** > **Workload Domains**, click the workload domain name, and then click the **Certificates** tab. The error appears in the **Status** column.
2. Review the information to make sure it is accurate and then click **Trust Certificate**.

Remove Old or Unused Certificates from SDDC Manager

Old or unused certificates are stored in a trust store in SDDC Manager. You can delete old certificates using the VMware Cloud Foundation API.

See [Delete Trusted Certificate](#) in the *VMware Cloud Foundation API Reference Guide* for more information.

1. Log in to the SDDC Manager UI as a user with the ADMIN role.
For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).
2. In the navigation pane, click **Developer Center** > **API Explorer**.
3. Browse to and expand **API Categories** > **Trusted Certificates**.
4. Expand **GET /v1/sddc-manager/trusted-certificates** and click **EXECUTE**.
5. In the Response, click `TrustedCertificate` and copy the alias for the certificate you want to remove.
6. Expand **DELETE /v1/sddc-manager/trusted-certificates/{alias}**, enter the alias, and click **EXECUTE**.

Try it out

Parameter	Value	Type	Description / Data Type
alias (required)	<input type="text"/>	path	Certificate Alias Data Type: string

EXECUTE

COPY RESPONSE

DOWNLOAD

Managing License Keys in VMware Cloud Foundation

You can add component license keys in the SDDC Manager UI or add a solution license key in vSphere Client.

Starting with VMware Cloud Foundation 5.1.1, you can license VMware Cloud Foundation components using a solution license key or individual component license keys.

NOTE

VMware Cloud Foundation 5.1.1 supports a combination of solution and component license keys. For example, `Workload Domain 1` can use component license keys and `Workload Domain 2` can use the solution license key.

For more information about the VCF solution license key, VMware vSphere 8 Enterprise Plus for VCF, see <https://knowledge.broadcom.com/external/article?articleNumber=319282>.

SDDC Manager does not manage the solution license key. If you are using a solution license key, VMware Cloud Foundation components are deployed in evaluation mode and then you use the vSphere Client to add and assign the

solution key. See [Managing vSphere Licenses](#) for information about using a solution license key for VMware ESXi and vCenter Server. If you are using a solution license key, you must also add a separate VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

NOTE

VMware vCenter Server, VMware NSX, VMware Aria Suite components, and VMware HCX are all licensed when you assign a solution license key to a vCenter Server.

Use the SDDC Manager UI to manage component license keys. If you entered component license keys in the deployment parameter workbook that you used to create the management domain, those component license keys appear in the Licensing screen of the SDDC Manager UI. You can add additional component license keys to support your requirements. You must have adequate license units available before you create a VI workload domain, add a host to a vSphere cluster, or add a vSphere cluster to a workload domain. Add the necessary component license keys before you begin any of these tasks.

Add a Component License Key in the SDDC Manager UI

You can use the SDDC Manager UI to add component license keys to the SDDC Manager inventory.

SDDC Manager does not manage solution license keys. See [Managing License Keys in VMware Cloud Foundation](#) for more information about solution license keys.

1. In the navigation pane, click **Administration** > **Licensing**.
2. Click **+ License Key**.



3. Select a product from the drop-down menu.
4. Enter the license key.
5. Enter a description for the license.
A description can help in identifying the license.
6. Click **Add**.

If you want to replace an existing license with a newly added license, you must add and assign the new license in the management UI (for example, vSphere Client or NSX Manager) of the component whose license you are replacing.

Edit a Component License Key Description in the SDDC Manager UI

If you have multiple component license keys for a product, the description can help in identifying the license key. For example, you may want to use one license key for high-performance workload domains and the other license key for regular workload domains.

1. In the navigation pane, click **Administration** > **Licensing**.
2. Click the vertical ellipsis (three dots) next to the license key and click **Edit Description**.

Edit Description

Remove

3. On the **Edit License Key Description** dialog, edit the description and click **Save**.

Delete a Component License Key in the SDDC Manager UI

Deleting a component license key removes it from the SDDC Manager inventory. If the license key has been applied to any workload domain, host, or vSphere cluster, it is not removed from them, but it cannot be applied to new workload domains, hosts, or vSphere clusters.

1. In the navigation pane, click **Administration** > **Licensing**.
2. Click the vertical ellipsis (three dots) next to the license key you want to delete and click **Remove**.

Edit Description

Remove

3. In the **Remove License key** dialog, click **Remove**.

The component license key is removed from the SDDC Manager inventory

Update Component License Keys for Workload Domain Components

You can use the SDDC Manager UI to update the license keys for components whose license keys have expired, are expiring, or are incompatible with upgraded components.

The new component license key(s) must already be added to the SDDC Manager inventory. See [Add a Component License Key in the](#) .

You can update component license keys for:

- vCenter Server
- VMware NSX
- VMware vSAN
- ESXi

Updates are specific to the selected workload domain. If you want to update component license keys for multiple workload domains, you must update each workload domain separately.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. Click a workload domain name in the **Domain** column.
3. Select **Actions** > **Update Licenses**.
4. Read the overview and click **Next**.
5. Select one or more products to update and click **Next**.

Update Licenses

> ✓ Overview Overview of updating licenses for a workload domain

2. Product Selection vCenter

Select product(s) within this workload domain that require the license to be updated.

Select products

<input type="checkbox"/>	Product	License key status ?	Available compatible license keys
<input checked="" type="checkbox"/>	vCenter	✓ Active	Yes
<input type="checkbox"/>	NSX	✓ Active	Yes
<input type="checkbox"/>	vSAN	✓ Active	Yes
<input type="checkbox"/>	ESXi	✓ Active	Yes

[NEXT](#)

3. vCenter License Apply license to vCenter

4. Review Review applied licenses

6. Select a component license key for each product.

For VMware vSAN and ESXi, you must select the clusters that you want to update with new license keys.

7. Review the new component license keys and click **Submit**.

Prepare ESXi Hosts for VMware Cloud Foundation

Before you can create a new VI workload domain, add a cluster to a workload domain, or add hosts to a cluster, you must prepare the ESXi hosts.

VI workload domains require a minimum of three ESXi hosts.

NOTE

If you are preparing ESXi hosts for a vSphere cluster using NFS, VMFS on FC, or vVols as principal storage, and the hosts will be added to a VI workload domain using vSphere Lifecycle Manager images as the update method, then only two hosts are required.

To use vSAN Express Storage Architecture (ESA), your hosts must be ESA-compatible.

TIP

See the [vSAN ESA VCG](#) for information about compatible hardware.

Preparing the ESXi hosts involves installing the correct version of ESXi and performing some basic configuration tasks. For the supported ESXi version, see the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes*.

IMPORTANT

If you are preparing hosts for a VI workload domain where the ESXi hosts have been async patched to a later version of ESXi than the version listed in the BOM, the new hosts must use the later version of ESXi.

Create a Custom ISO Image for ESXi

When your environment requires a custom ISO file for ESXi, you can create one using VMware PowerCLI or vSphere Lifecycle Manager.

Download the zip files for the following:

- ESXi patch for the ESXi version specified in the VMware Cloud Foundation BOM or in the list of supported async patches in [KB 88287](#). You can download patches from the Broadcom Support Portal.

NOTE

If you are preparing hosts for a VI workload domain where the ESXi hosts have been async patched to a later version of ESXi than the version listed in the BOM, the new hosts must use the later version of ESXi.

- OEM add-on for ESXi from the Broadcom Support Portal. If the ESXi version specified in the BOM is not available in the **Select Version** drop-down menu, contact your vendor to determine which OEM add-on version to use.

You might need to create a custom ISO image for ESXi in the following situations:

- The ESXi version specified in the VMware Cloud Foundation BOM does not have an associated ISO file on the Broadcom Support Portal. This can be the case for ESXi patch releases.
- You need an async patch version of ESXi.
- You need a vendor-specific (OEM) ISO file.

Create a Custom ESXi ISO Image Using VMware PowerCLI

You can use VMware Power CLI to create a custom ISO.

VMware PowerCLI 12.0 or later.

1. Gather the required information for the software spec that is used to create the custom ISO.
 - a) In VMware PowerCLI, use the [Get-DepotBaseImages](#) cmdlet to get the base image version from the zip file for the ESXi patch that you downloaded from the patches portal.

For example:

```
Get-DepotBaseImages "c:\temp\VMware-ESXi-7.0U1d-17551050-depot.zip"
```

- b) Use the [Get-DepotAddons](#) cmdlet to get the add-on name and version from the zip file for the OEM add-on for ESXi that you downloaded from the Broadcom Support Portal. (if applicable)

For example:

```
Get-DepotAddons "c:\temp\HPE-701.0.0.10.6.5.12-Jan2021-Synergy-Addon-depot.zip"
```

2. Create the software spec using the information you gathered in step 1.

The software spec is a JSON file that contains information about the ESXi version and vendor add-on (if applicable). For example:

```
{
  "add_on": {
    "name": "HPE-Custom-Syn-AddOn",
    "version": "701.0.0.10.6.5-12"
  },
  "base_image": {
    "version": "7.0.1-0.30.17551050"
  },
}
```

```

"components": null,
"hardware_support": null,
"solutions": null
}

```

3. In VMware PowerCLI, use the `New-IsoImage` cmdlet to generate a custom ISO.

For example:

```

New-IsoImage -SoftwareSpec "c:\temp\HPE-70Uld-custom.JSON" -Depots "c:\temp\VMware-ESXi-7.0Uld-17551050-depot.zip" , "c:\temp\HPE-701.0.0.10.6.5.12-Jan2021-Synergy-Addon-depot.zip" -Destination "c:\temp\HPE-70Uld-custom.iso"

```

Provide the path to the software spec you created in step 2.

The depot(s) include the path to the zip files for the supported ESXi version and vendor add-on.

The destination include the path and file name for the custom ISO file.

For more information about the `New-IsoImage` cmdlet, see <https://developer.broadcom.com/powercli/latest/vmware.imagebuilder/commands/new-isoimage>.

Create a Custom ESXi ISO Image Using vSphere Lifecycle Manager

If you have access to a vCenter Server 7.0 environment, you can use vSphere Lifecycle Manager to create and export a custom ISO.

Import the ESXi patch and vendor add-on (if applicable) zip files to the vSphere Lifecycle Manager depot. See [Import Updates to the vSphere Lifecycle Manager Depot](#).

1. Log in to vCenter Server using the vSphere Client.
2. Create a new temporary cluster, selecting the **Manage all hosts in the cluster with a single image** check box.
3. Select the ESXi version and vendor add-on (optional) and click **OK**.
4. Export the vSphere Lifecycle Manager image as an ISO.

See [Export an Image](#).

5. Delete the temporary cluster.

Install ESXi Interactively and Configure Hosts for VMware Cloud Foundation

You can interactively install ESXi on all the hosts that will form the first cluster in the management domain, then you configure the management network, DNS, and NTP services. You can use the same process to add more hosts to the management domain later, or to install and configure hosts for VI workload domains.

- Download the ESXi ISO from the Broadcom Support Portal. For the supported ESXi versions, see the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes* and the list of supported async patches in [KB 88287](#). If the required version of ESXi does not have an ISO available on the Broadcom Support Portal, you can create one. See [Create a Custom ISO Image for ESXi](#).

NOTE

If you are preparing hosts for a VI workload domain where the ESXi hosts have been async patched to a later version of ESXi than the version listed in the BOM, the new hosts must use the later version of ESXi.

- Make sure that you have a host machine for SDDC access. You use this host to connect to the data center and perform configuration steps.
- Verify that you have the completed *Planning and Preparation Workbook*.
- Verify the Prerequisite Checklist sheet in the *Planning and Preparation Workbook*.

ESXi 8.0 Update 3 and later support installing two data processing units (DPUs) for use with VMware Cloud Foundation 5.2 or later.

You can utilize the two DPUs in Active/Standby mode to provide high availability. Such configuration provides redundancy in the event one of the DPUs fails. In the high availability configuration, both DPUs are assigned to the same NSX-backed vSphere Distributed Switch. For example, DPU-1 is attached to vmnic0 and vmnic1 of the vSphere Distributed Switch and DPU-2 is attached to vmnic2 and vmnic3 of the same vSphere Distributed Switch.

You can also utilize the two DPUs as independent devices to increase offload capacity per ESXi host. Each DPU is attached to a separate vSphere Distributed Switch and you have no failover between DPUs in such configuration.

Install ESXi on VMware Cloud Foundation Hosts Using the ISO

Install ESXi on all hosts in the first cluster in the management domain interactively. You can use the same process to install ESXi on additional hosts for the management domain, or on hosts for a VI workload domain.

Repeat this procedure for each host in the cluster that you want to add to a workload domain.

1. Mount the ESXi ISO on the host and restart the machine.
2. Set the BIOS or UEFI to boot from the mounted ISO.

NOTE

If your system has supported data processing units (DPUs), you can only use UEFI to install and boot ESXi on the DPUs.

See your hardware vendor documentation for information on changing boot order.

3. On the welcome screen, press **Enter** to continue.
4. Accept the End User License Agreement by pressing **Enter**.

Starting with ESXi 8.0 Update 3, after the scanning for available devices completes, if your system has DPUs, you see them automatically listed with their respective PCI slots. You no longer select a slot. The DPU devices must be identical: same vendor, same hardware version and same firmware

5. On the **Select a Disk to Install or Upgrade** screen, select the drive on which to install ESXi on and press **Enter**.
6. Select the keyboard type for the host.

You can change the keyboard type after installation in the direct console.

7. Enter the root password for the host.
8. In the **Confirm Install** screen, if you have DPUs, you see each listed on a separate row. Press **F11** to confirm the start of the installation.

Starting with ESXi 8.0 Update 3, if your systems has DPUs, you see a single progress bar for the ESXi and DPU installation, with dynamic updates to the label showing what stage of the installer is being run.

9. On the **Installation Complete** screen, press **Enter** to reboot the host.
10. Set the first boot device to be the drive on which you installed ESXi.
11. Repeat this procedure for all remaining hosts.

Configure the Network on VMware Cloud Foundation Hosts

After the initial boot, use the ESXi Direct Console User Interface (DCUI) for host network configuration and administrative access.

Perform the following tasks to configure the host network settings:

- Configure the network adapter (vmk0) and VLAN ID for the Management Network.
- Configure the IP address, subnet mask, gateway, DNS server, and FQDN for the ESXi host.

Repeat this procedure for all the hosts that you are adding to the workload domain.

1. Open the DCUI of the ESXi host.
 - a) Open a console window to the host.
 - b) Press F2 to enter the DCUI.
 - c) Log in by using the `esxi_root_user_password`.
2. Configure the network.
 - a) Select **Configure Management Network** and press Enter.
 - b) Select **VLAN (Optional)** and press Enter.
 - c) Enter the VLAN ID for the Management Network and press Enter.
 - d) Select **IPv4 Configuration** and press Enter.
 - e) Select **Set static IPv4 address and network configuration** and press the Space bar.
 - f) Enter the IPv4 Address, Subnet Mask and Default Gateway and press Enter.
 - g) Select **DNS Configuration** and press Enter.
 - h) Select **Use the following DNS Server address and hostname** and press the Space bar.
 - i) Enter the Primary DNS Server, Alternate DNS Server and Hostname (FQDN) and press Enter.
 - j) Select **Custom DNS Suffixes** and press Enter.
 - k) Ensure that there are no suffixes listed and press Enter.
3. Press Escape to exit and press Y to confirm the changes.
4. Repeat this procedure for all remaining hosts.

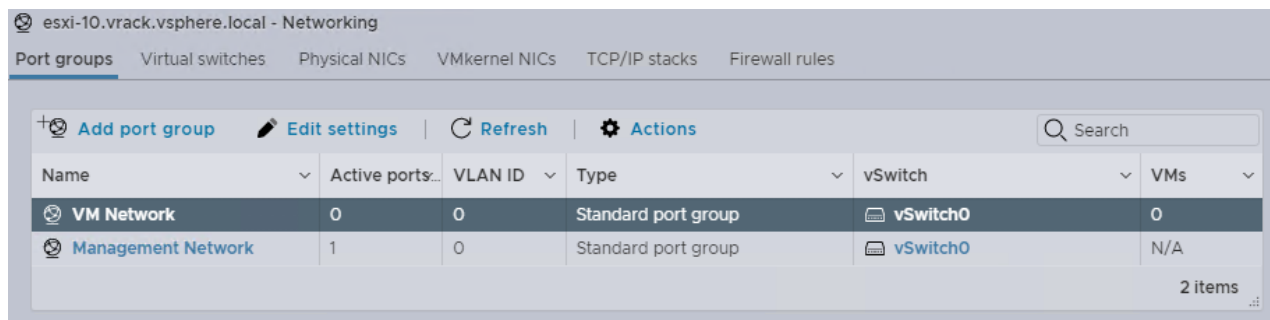
Configure the Virtual Machine Network Port Group on VMware Cloud Foundation Hosts

You perform configuration of the Virtual Machine Network port group for each ESXi host by using the VMware Host Client.

You configure the VLAN ID of the VM Network port group on the vSphere Standard Switch. This configuration provides connectivity to the Management network to allow communication to the vCenter Server Appliance during the automated deployment.

Repeat this procedure for each host that you are adding to the workload domain.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. Click **OK** to join the Customer Experience Improvement Program.
3. Configure a VLAN for the VM Network port group.
 - a) In the navigation pane, click **Networking**.
 - b) Click the **Port groups** tab, select the **VM network** port group, and click **Edit Settings**.



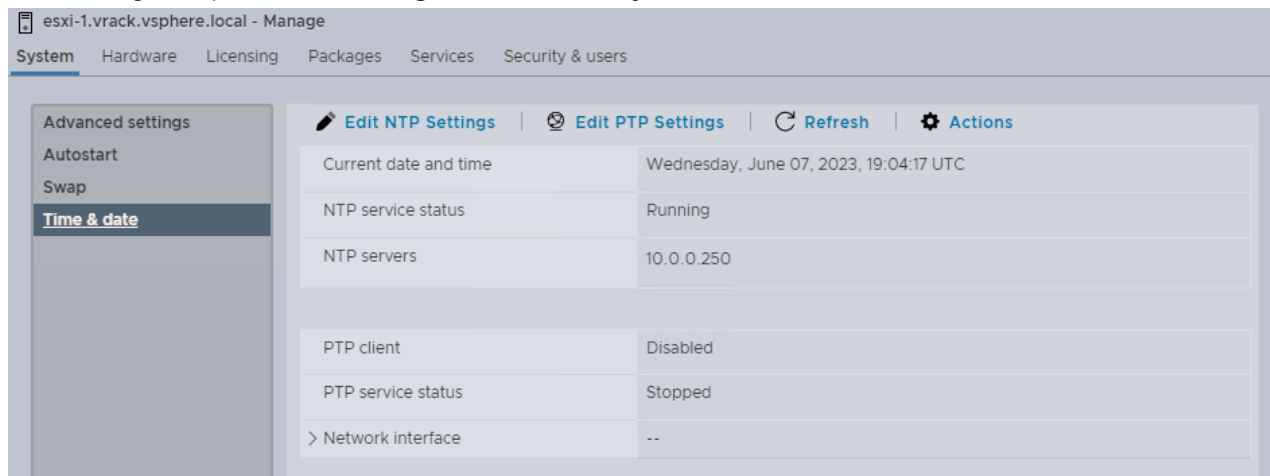
- c) On the **Edit port group - VM network** page, enter the Management Network VLAN ID, and click **Save**.
4. Repeat this procedure for all remaining hosts.

Configure NTP on VMware Cloud Foundation Hosts

Complete the initial configuration of all ESXi hosts by configuring the NTP service to avoid time synchronization issues in the SDDC.

Repeat this procedure for all the hosts that you are adding to the workload domain.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. Configure and start the NTP service.
 - a) In the navigation pane, click **Manage**, and click the **System** tab.



- b) Click **Time & date** and click **Edit NTP Settings**.
- c) On the **Edit NTP Settings** page, select the **Use Network Time Protocol (enable NTP client)** radio button, and change the NTP service startup policy to **Start and stop with host**.
- d) In the **NTP servers** text box, enter the NTP Server FQDN or IP Address, and click **Save**.
- e) To start the service, click **Actions**, select **NTP service**, and click **Start**.
3. Repeat this procedure for all remaining hosts.

Regenerate the Self-Signed Certificate on All Hosts

Once you have configured the ESXi hosts' identity by providing a hostname you must regenerate the self-signed certificate to ensure the correct common name is defined.

During the installation of ESXi, the installer generates a self-signed certificate for each ESXi host but the process is performed prior to the ESXi identity being configured. This means all ESXi hosts have a common name in their self-signed certificate of `localhost.localdomain`. All communication between VMware Cloud Builder and the ESXi hosts is

performed securely over HTTPS and as a result it validates the identify when making a connection by comparing the common name of the certificate against the FQDN provided within the VMware Cloud Builder configuration file. To ensure that the connection attempts and validation does not fail, you must manually regenerate the self-signed certificate after hostname has been configured.

NOTE

VMware Cloud Foundation supports the use of signed certificates. If your organization's security policy mandates that all ESXi hosts must be configured with a CA-signed certificate, see [Configure ESXi Hosts with Signed Certificates](#).

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the **Actions** menu, click **Services > Enable Secure Shell (SSH)**.
3. Log in to the ESXi host using an SSH client such as Putty.
4. Regenerate the self-signed certificate by executing the following command:

```
/sbin/generate-certificates
```

5. Reboot the host.
6. Log back in to the VMware Host Client and click **Services > Disable Secure Shell (SSH)** from the **Actions** menu.
7. Repeat this procedure for all remaining hosts.

Configure ESXi Hosts with Signed Certificates

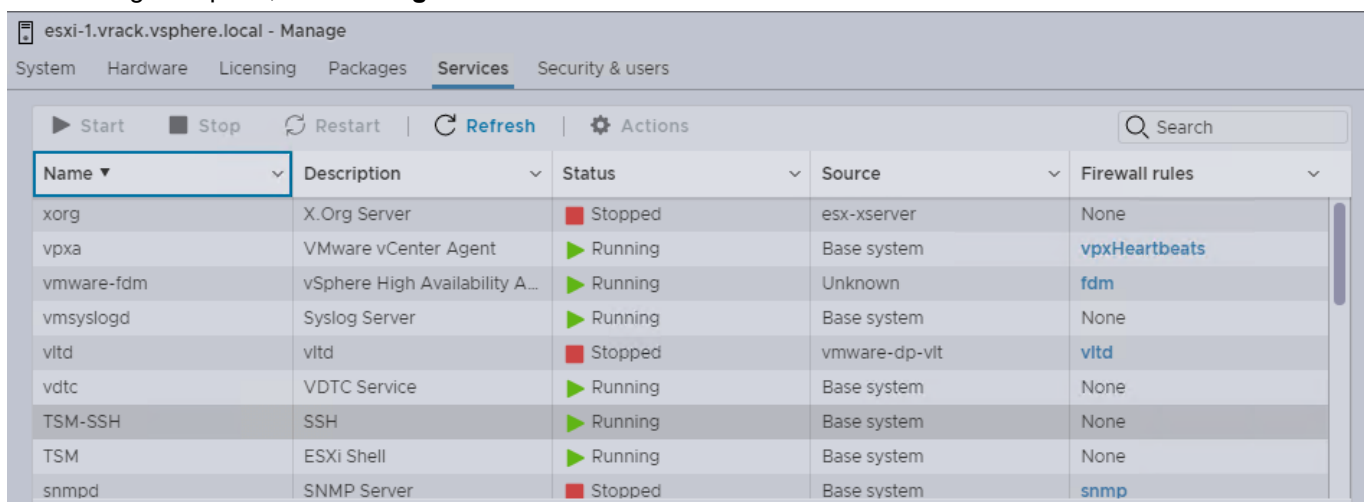
If corporate policy requires that you use external CA-signed certificates instead of VMCA-signed certificates for ESXi hosts, you can manually add external certificates to the hosts.

External CA-signed certificate and key are available.

When you install ESXi software on a server to create an ESXi host, the host initially has an autogenerated certificate. By default, when the host is added to a vCenter Server system during bring-up of the management domain or other operations involving hosts (for example, host commissioning, VI workload domain creation, and so on), the autogenerated certificate is replaced with a certificate that is signed by the VMware Certificate Authority (VMCA).

When you use external certificates during bring-up, they are not replaced by VMCA-signed certificates. Once you perform bring-up with external certificates for ESXi hosts, all future hosts added to VMware Cloud Foundation must also use external certificates.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the navigation pane, click **Manage** and click the **Services** tab.



3. Select the **TSM-SSH** service and click **Start** if not started.

4. Log in to the ESXi Shell for the first host, either directly from the DCUI or from an SSH client, as a user with administrator privileges.
5. In the directory `/etc/vmware/ssl`, rename the existing certificates using the following commands:


```
mv rui.crt orig.rui.crt
mv rui.key orig.rui.key
```
6. Copy the external certificate and key that you want to use to `/etc/vmware/ssl`.
7. Rename the external certificate and key to `rui.crt` and `rui.key`.
8. Restart the host management agents by running the following commands:


```
/etc/init.d/hostd restart
/etc/init.d/vpxa restart
```
9. In the VMware Host Client, select the **TSM-SSH** service for the ESXi host and click **Stop**.
10. Repeat for all the ESXi hosts that you are adding to VMware Cloud Foundation.

Managing ESXi Hosts in VMware Cloud Foundation

To add ESXi hosts to the SDDC Manager inventory, you must first create a network pool or expand the default network pool created during bring-up. If you want to isolate VMkernel traffic (management, vSAN, vMotion, or overlay) across multiple physical NICs (pNICs) on the hosts, you must authorize the additional pNICs before adding the hosts to VMware Cloud Foundation.

For information on network pools, see [Network Pool Management](#).

During the commissioning process, the hosts are associated with a network pool and added to the SDDC Manager inventory. Newly commissioned hosts appear in the SDDC Manager inventory with a host state of UNASSIGNED, and can now be added to either the management domain or a VI workload domain, depending on their associated network pool. When a host is added to a workload domain, an IP address from the network pool's IP inclusion range is assigned to it.

See [VMware Configuration Maximums](#) for information about the maximum number of hosts per SDDC Manager instance.

Network Pool Management

When you create a VI workload domain or add a host or vSphere cluster to an existing workload domain (management or VI), you do not need to enter IP addresses manually. Network pools automatically assign static IP addresses to vSAN, NFS, iSCSI, and vMotion VMkernel ports.

A network pool is a collection of subnets within an layer-2 network domain. Depending on the storage option, it includes information about subnets reserved for the vMotion and vSAN, NFS, or iSCSI networks that are required for adding a host to the SDDC Manager inventory.

Table 180: Information Required for a Network Pool

Storage Being Used	Required Networks in Network Pool
vSAN	vMotion and vSAN
vSAN Max	vMotion and vSAN
NFS	vMotion and NFS
vSAN and NFS	vMotion, vSAN, and NFS
VMFS on FC	vMotion only or vMotion and NFS
vVols on FC	vMotion only or vMotion and NFS

Table continued on next page

Continued from previous page

Storage Being Used	Required Networks in Network Pool
vVols on iSCSI	vMotion and iSCSI
vVols on NFS	vMotion and NFS

The network pool also contains a range of IP addresses, called an inclusion range. IP addresses from the inclusion ranges are assigned to the VMkernel ports on the host. The use of inclusion ranges allows you to limit the IP addresses that are consumed from a given subnet. You can add more inclusion ranges to expand the use of the provided subnet.

A default network pool is created during bring-up. This network pool is automatically associated with the management domain. Network information for this network pool is based on the deployment parameter workbook you provided during bring-up. This network pool contains vMotion and vSAN networks only - an NFS network is not supported in this network pool. If the vSAN and vMotion networks in your management domain are in the same layer-2 network domain, you can expand the default network pool. You can also expand the default network pool if you expand the management domain by adding a host.

To create a VI workload domain with hosts in a different layer-2 network domain than the management domain, you must create a new network pool. Also, if you want to use external NFS or VMFS on FC storage, you must create a new network pool. A network pool can contain both vSAN and NFS networks.

You can also create a workload domain with multiple vSphere clusters, each with its own network pool. You can have multiple vSphere clusters within a workload domain to provide a separate fail over domain (a VM only fails over between hosts in a cluster). Multiple vSphere clusters also provide isolation for security reasons and are also useful for grouping servers of a particular type of configuration together. Multiple vSphere clusters can also be used to handle the growth. The original servers used in the default vSphere cluster can get outdated at some point. Newer server models can then be added to a new cluster in the workload domain and workloads can be migrated at a leisurely pace.

Size a Network Pool

Properly sizing a network pool is critical to prevent future issues in the environment due to insufficient IP addresses. Care must be taken when defining the subnets for a network pool as the subnet cannot be changed after it is deployed. The scope of IP addresses used from the defined subnet can be limited by the definition of one or more inclusion ranges. Thus, it is recommended that you begin with defining a larger subnet than what is initially required and utilize the inclusion ranges to limit use. This will provide you the capability to grow with demand as needed.

You begin sizing a network pool by determining the number of hosts that you will have in each vSphere cluster. A VI workload domain must contain a minimum of one vSphere cluster, with a minimum number of three hosts. The management workload domain requires a minimum of four hosts, which allows for an additional level of availability for the critical infrastructure components. A vSphere cluster can be expanded to the maximum number of hosts supported by vCenter, which is currently 64 hosts.

NOTE

If vSphere cluster is using NFS, VMFS on FC, or vVols as principal storage, and the cluster will be added to a VI workload domain using vSphere Lifecycle Manager images as the update method, then only two hosts are required.

Allocate a minimum of one IP address per host plus enough additional IP addresses to account for growth and expansion of the environment. Ensure that the subnet defined provides enough unused IP addresses and that appropriate inclusion ranges are defined. Note that some of the IP addresses within the subnet will be used for other purposes, such as defining the gateway address, firewalls, or other entities. Use care not to conflict with these addresses.

Here are some important considerations for determining the size of your network pool:

- Type of network architecture
- Physical switch details
 - Are they managed or non-managed?

Do they support layer-3? (this may require a license)

Number of ports

- Where the network switches are placed (at the top of the rack or at the end of a row)
- Where the default gateway is created
- Number of hosts that can be placed in each rack or layer-2 network domain
- Number of hosts required in a vSphere cluster
- Whether the network switches will be shared with non- VMware Cloud Foundation hosts
- Number of workload domains you plan on creating

View Network Pool Details

You can view the network details for a network pool as well as the total number of used and available IP addresses.

1. In the navigation pane, click **Administration** > **Network Settings**.
2. Click the arrow to the left of the pool name.

Network Pool Name	Hosts Per Pool																				
bringup-networkpool	13																				
<table border="1"> <thead> <tr> <th colspan="2">vMotion</th> <th colspan="2">vSAN</th> </tr> </thead> <tbody> <tr> <td>Network</td> <td>10.0.8.0</td> <td>Network</td> <td>10.0.4.0</td> </tr> <tr> <td>Subnet Mask</td> <td>255.255.255.0</td> <td>Subnet Mask</td> <td>255.255.255.0</td> </tr> <tr> <td>Free IPs</td> <td>36</td> <td>Free IPs</td> <td>36</td> </tr> <tr> <td>Used IPs</td> <td>12</td> <td>Used IPs</td> <td>12</td> </tr> </tbody> </table>		vMotion		vSAN		Network	10.0.8.0	Network	10.0.4.0	Subnet Mask	255.255.255.0	Subnet Mask	255.255.255.0	Free IPs	36	Free IPs	36	Used IPs	12	Used IPs	12
vMotion		vSAN																			
Network	10.0.8.0	Network	10.0.4.0																		
Subnet Mask	255.255.255.0	Subnet Mask	255.255.255.0																		
Free IPs	36	Free IPs	36																		
Used IPs	12	Used IPs	12																		
my-nfs-vsant-networkpool	0																				

2 Network Pools

A summary view of the network pools details is displayed.

3. Click the name of a network pool.
A detailed view of the network pools details is displayed.

Create a Network Pool

A network pool must include the vMotion network information. Depending on the type of storage you are using, you must provide network information for vSAN, NFS, and iSCSI.

The subnet in a network pool cannot overlap the subnet of another pool.

1. In the navigation pane, click **Administration** > **Network Settings**.
2. Click **Create Network Pool**.

Create Network Pool

Ensure that all required networks are selected based on their usage for workload domains.

Network Pool Name

Network Type ⓘ vSAN NFS iSCSI vMotion

vMotion Network Information

VLAN ID ⓘ	<input type="text" value="VLAN ID"/>
MTU ⓘ	<input type="text" value="9000"/>
Network ⓘ	<input type="text" value="XXX.XXX.XXX.XXX"/>
Subnet Mask ⓘ	<input type="text" value="XXX.XXX.XXX.XXX"/>
Default Gateway ⓘ	<input type="text" value="XXX.XXX.XXX.XXX"/>

Included IP Address Ranges

Once a network pool has been created, you are not able to edit or remove IP ranges from that pool.

To

3. Enter a name for the network pool.
4. Select the network type(s).
 - For vSAN and NFS storage, you can include both vSAN and NFS network information in the same network pool or create separate network pools for vSAN and NFS.
 - For vSAN Max storage, select vSAN and vMotion.
 - For VMFS on FC storage, select vMotion only or vMotion and NFS.
 - For vVols on FC storage, select vMotion only or vMotion and NFS.
 - For vVols on iSCSI storage, select vMotion and iSCSI.
 - For vVols on NFS storage, select vMotion and NFS.
5. Provide the following network information for the selected network type(s).


Option	Description
VLAN ID	Enter a VLAN ID between 1 and 4094.
MTU	Enter an MTU between 1500 and 9000.
Network	Enter a subnet IP address.
Subnet Mask	Enter the subnet mask.
Default Gateway	Enter the default gateway.
Included IP Ranges	<p>Enter an IP address range from which an IP address can be assigned to hosts that are associated with this network pool. The IP address range must be from within the specified subnet. You cannot include the IP address of the default gateway in the IP address range. You can enter multiple IP address ranges.</p> <p>NOTE Ensure that you have entered the correct IP address range. IP ranges cannot be edited after the network pool is created.</p>

6. Click **Save**.

Add or Remove a Network Pool IP Address Range

You can edit a network pool to add a new IP address range or to remove an existing IP address range.

1. In the navigation pane, click **Administration > Network Settings**.
2. Click the vertical ellipsis (three dots) in the network pool row you want to edit and then click **Edit Network Pool IP Range**.



3. To add a new IP address range, enter the start and end IP addresses and click **Add**.
4. To remove an existing IP address range, click **Remove**.

NOTE

You cannot remove an existing IP address range if it is the only range that exists within the network pool or if it is in use.

Rename a Network Pool

You can rename an existing network pool.

1. In the navigation pane, click **Administration > Network Settings**.
2. Click the vertical ellipsis (three dots) in the network pool row you want to rename and then click **Rename Network Pool**.

3. Enter a new name for the network pool.
4. Click **Rename**.

Delete a Network Pool

You can delete a network pool if none of the hosts with an IP address from the pool belong to a workload domain. The default pool created during bring-up cannot be deleted.

Ensure that the hosts in the network pool are not assigned to a workload domain. To verify this, navigate to **Administration > Network Settings** and confirm that the **Used IPs** for the network pool is 0.

1. In the navigation pane, click **Administration > Network Settings**.
2. Click the vertical ellipsis (three dots) in the network pool row that you want to delete and click **Delete**.

Rename Network Pool
 Edit Network Pool IP Range
 Delete

3. Click **Delete**.

View Host Inventory

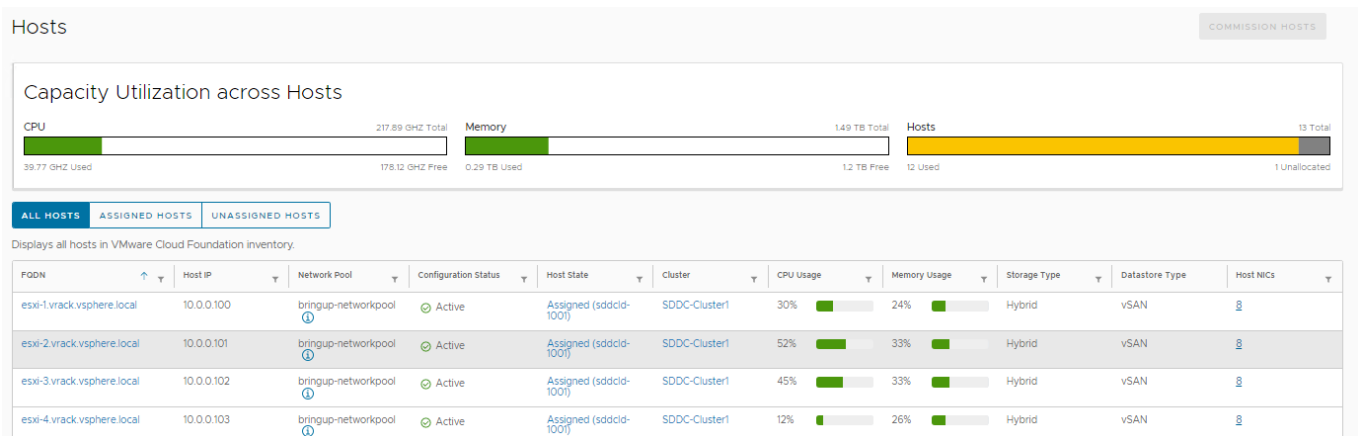
The Hosts page displays details about all the hosts in the SDDC Manager inventory, including CPU utilization and memory usage across all hosts, as well as the total number of hosts used and unallocated.

For each host, the Hosts page displays the following information:

- FQDN
- Host IP address
- The network pool to which the host belongs
- Configuration status
- Host state and workload domain
- vSphere cluster
- CPU and memory usage
- Storage type

The Hosts page also provides controls for commissioning hosts.

1. In the navigation pane, click **Inventory > Hosts**.



2. Navigate directly to pages related to a specific host.

For example:

- To jump to the details page for the domain to which a listed host belongs, click the domain name link in the Host State column. For information about viewing workload domains, see [View Workload Domain Details](#).
- To jump to the details page for the domain cluster to which a listed host belongs, click the cluster name in the Cluster column. For information about clusters, see [Expand a Workload Domain](#).
- To quickly view network assignment details for a specific host, click the info icon next to the value in the Network Pool column.

3. To view the details of a specific host, click the FQDN in the list.

The host details page appears, displaying the following information:

- A chart showing total and used CPU capacity.
- A chart showing total and used memory capacity.
- A summary of the networks (vSAN, vMotion, and Management) to which the host belongs and its IP address on those networks.
- The manufacturer and model of the host.
- Storage information for capacity and cache tiers.
- ESXi version.

NOTE

Below the page title, the host details page also provides quick links to the network pool and the vSphere cluster to which the host belongs.

4. To view additional configuration details for the ESXi host, click **Actions** › **Open in VMware Host Client**.

Commission Hosts

Adding hosts to the SDDC Manager inventory is called commissioning. You can add hosts individually or use a JSON template to commission multiple hosts at once. You can also run multiple commission hosts tasks at the same time.

Ensure that each host you are commissioning meets the following criteria:

- Hosts for vSAN-based workload domains are vSAN-compliant and certified on the VMware Hardware Compatibility Guide.
- Hosts for NFS-based workload domains are certified on the VMware Hardware Compatibility Guide.
- Hosts for VMFS on FC-based workload domains are certified on the VMware Compatibility Guide. In addition, the hosts must have supported FC cards (Host Bus Adapters) and drivers installed and configured. For compatible FC cards, see the VMware Compatibility Guide.

IMPORTANT

The VMFS datastore must be mounted on all the hosts before commissioning.

- Hosts for vVols-based workload domains are certified on the VMware Hardware Compatibility Guide.
 - For vVols on FC-based workload domains, ensure that all ESXi hosts have access to the FC array before launching the workflow.
 - For vVols on NFS-based workload domains, ensure that all ESXi hosts must be able to reach the NFS server from the NFS network assigned in the IP pool.
 - For vVols on iSCSI-based workload domains, ensure that the iSCSI software initiator must be enabled on each ESXi host and the VASA provider URL must be listed as the dynamic target.
- For hosts with a DPU device, enable SR-IOV in the BIOS and in the vSphere Client (if required by your DPU vendor).
- Two NIC ports with a minimum 10 Gbps speed. One port must be free, and the other port must be configured on a standard switch. This switch must be restricted to the management port group.
- Host has the drivers and firmware versions specified in the VMware Hardware Compatibility Guide.
- A supported version of ESXi is installed on the host. See the *VMware Cloud Foundation Release Notes* for information about supported versions.

- DNS is configured for forward and reverse lookup and FQDN.
- Host name must be same as the FQDN.
- Self-signed certificate regenerated based on FQDN of host. See [Regenerate the Self-Signed Certificate on All Hosts](#).
- Management IP address is configured on the first NIC port.
- Host has a standard switch and configured with 10 Gbps speed default uplinks can start with vmnic1, 2, 3, and so forth.
- Hardware health status is healthy without any errors.
- All disk partitions on HDD and SSD are deleted.
- Network pool must be created and available before host commissioning.
- Hosts for the vSAN-based workload domain must be associated with the vSAN enabled network pool.
- Hosts for the NFS-based workload domain are associated with the NFS enabled network pool.
- Host is configured with an appropriate gateway. The gateway must be part of the management subnet.

The hosts that you want to commission must meet a set of criteria. After you specify host details and select the network pool to associate a host with, VMware Cloud Foundation validates and commissions each host. Each host is added to the free pool and is available for VI workload domains and vSphere clusters.

- Hosts that use vSAN storage can only be used with vSAN-based workload domains.
 - Hosts that are commissioned for vSAN ESA can only be used for vSAN ESA clusters.
 - Hosts that are commissioned for vSAN OSA can only be used for vSAN OSA clusters.
 - Hosts that are commissioned for vSAN Max can only be used for vSAN Max clusters.
 - Hosts that are commissioned for vSAN Compute Clusters (VSAN_REMOTE) can only be used for vSAN compute clusters.
- Hosts that use NFS storage can only be used with NFS-based workload domains.
- Hosts that use VMFS on FC storage can only be used with VMFS on FC-based workload domains.
- Hosts that use vVols storage can only be used with vVols-based workload domains.

See [VMware Configuration Maximums](#) for information about the maximum number of hosts you can commission at one time and the maximum number of commission hosts tasks that you can run in parallel.

Ensure that a network pool supports the storage type you select for a host (vSAN, vSAN ESA, vSAN Max, NFS, VMFS on FC, vVols). See [Network Pool Management](#).

If you used hosts with certificates generated by an external CA for bring-up, hosts used for expanding workload domains (management or VI) or for creating VI workload domains must also use certificates generated by an external CA. See [Configure ESXi Hosts with Signed Certificates](#).

NOTE

When you add a host during host commissioning, SDDC Manager temporarily activates SSH on the host in order to retrieve the host key. After it retrieves the host key, SDDC Manager deactivates SSH on the host.

1. In the navigation pane, click **Inventory** > **Hosts**.
2. Click **Commission Hosts**.
3. On the **Checklist** page, review the prerequisites and confirm by clicking **Select All**.
4. Click **Proceed**.
5. Select whether you want to add hosts one at a time or import multiple hosts at once from a JSON file.

Option	Description
Add new	Enter the following information for the host you want to add: <ul style="list-style-type: none"> • Host FQDN The host FQDN must match the host name (including capitalization) provided during DNS configuration. • Storage Type (vSAN, NFS, VMFS on FC, or vVol) for principal storage <ul style="list-style-type: none"> – If you select vSAN as the storage type:

Table continued on next page

Continued from previous page

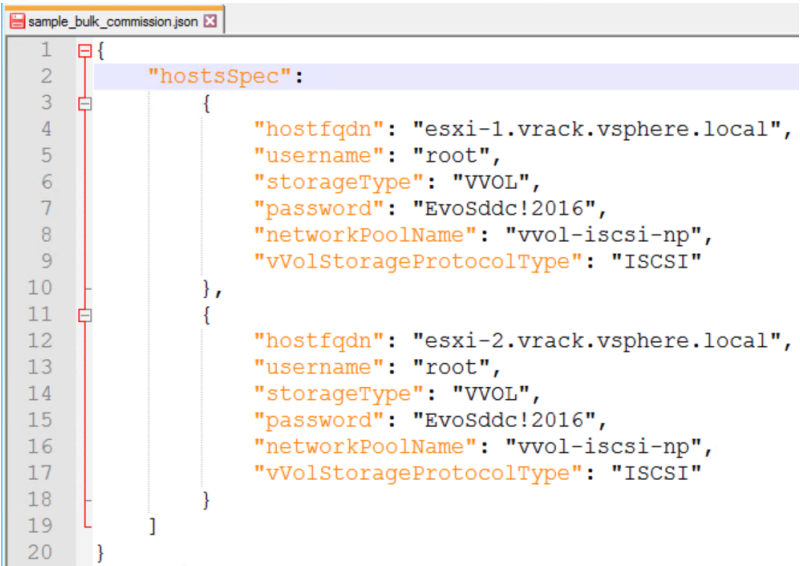
Option	Description
	<ul style="list-style-type: none"> • Select the vSAN Type (vSAN HCI, vSAN Max, or vSAN Compute Cluster). • For vSAN HCI choose whether or not to use the vSAN Express Storage Architecture. vSAN ESA is only supported in workload domains that use vLCM images. <ul style="list-style-type: none"> – If you select vVols as the storage type, specify the vVols storage protocol type (FC, iSCSI, or NFS). • Network Pool Name (select a network pool from the drop-down menu) • User Name and Password (root credentials for the ESXi host) <p>Click Add. You can now add more ESXi hosts or proceed to the next step.</p>
Import	<ol style="list-style-type: none"> 1. To download the JSON template, click the link. 2. Open the JSON template file and enter information about the hosts to add. <ul style="list-style-type: none"> – Host FQDN – Username and password (ESXi root credentials) – Storage type (VSAN, VSAN_REMOTE, VSAN_ESA, VSAN_MAX, NFS, VMFS_FC, VVOL) for principal storage <p>NOTE Use VSAN_REMOTE for a vSAN Compute Cluster.</p> <ul style="list-style-type: none"> – Network pool name <p>NOTE You only need to provide the vVols storage protocol type (FC, ISCSI, or NFS) when you are using VVOLS as the storage type</p>  <pre> 1 { 2 "hostsSpec": 3 { 4 "hostfqdn": "esxi-1.vrack.vsphere.local", 5 "username": "root", 6 "storageType": "VVOL", 7 "password": "EvoSddc!2016", 8 "networkPoolName": "vvol-iscsi-np", 9 "vVolStorageProtocolType": "ISCSI" 10 }, 11 { 12 "hostfqdn": "esxi-2.vrack.vsphere.local", 13 "username": "root", 14 "storageType": "VVOL", 15 "password": "EvoSddc!2016", 16 "networkPoolName": "vvol-iscsi-np", 17 "vVolStorageProtocolType": "ISCSI" 18 } 19] 20 } </pre> 3. To locate and select the JSON file containing host information, click Browse.

Table continued on next page

Continued from previous page

Option	Description
	4. Click Upload .

The host or hosts appear in the **Hosts Added** section.

6. Verify that the server fingerprint is correct for each host and then activate the **Confirm All Finger Prints** toggle.
7. Click **Validate All**.
VMware Cloud Foundation validates the host information you provided. Each host is marked as **Valid** or **Invalid**. For invalid hosts, you can correct the problem and validate again, or select the host and click **Remove** to proceed with commissioning the valid hosts.
8. Click **Next**.
9. Activate or deactivate **Skip failed hosts during commissioning**.
This setting is enabled by default. When enabled, hosts that fail are skipped and commissioning continues on the remaining hosts.
10. Click **Commission**.
The Hosts page appears, and the status of the commission task is displayed in the Tasks window.

NOTE

Multiple VMkernels are created to test the vMotion network, which may cause changes in the MAC addresses and IP address relations. If MAC address filtering is enabled on your physical infrastructure, this may cause issues such as vMotion network connectivity validation failure.

The commissioned hosts are added to the hosts table in the SDDC Manager inventory as unassigned hosts.

Decommission Hosts

Removing hosts from the SDDC Manager inventory is called decommissioning. If you want to re-use a host in a different workload domain, you must decommission, re-image, and commission the host before adding it to the workload domain.

The hosts that you want to decommission must not be assigned to a workload domain. If a host is assigned to a workload domain, you must remove it before you can decommission it. See [Remove a Host from a vSphere Cluster in a Workload Domain](#).

See [VMware Configuration Maximums](#) for information about the maximum number of hosts you can decommission at one time and the maximum number of decommission hosts tasks that you can run in parallel..

1. In the navigation pane, click **Inventory > Hosts**.
2. Click **Unassigned Hosts**.

ALL HOSTS ASSIGNED HOSTS UNASSIGNED HOSTS

Displays only decommission eligible hosts that are not assigned to any workload domains.

DECOMMISSION SELECTED HOSTS

<input type="checkbox"/>	FQDN	Host IP	Network Pool	Configuration Status	Host State	Cluster	CPU Usage	Memory Usage	Storage Type	Datastore Type	Host NICs
<input checked="" type="checkbox"/>	esxi-10.vrack.vsphere.local	10.0.0.109	bringup-networkpool	Active	Unassigned	-	1%	7%	Hybrid	-	2
<input type="checkbox"/>	esxi-11.vrack.vsphere.local	10.0.0.110	bringup-networkpool	Active	Unassigned	-	1%	7%	Hybrid	-	2
<input type="checkbox"/>	esxi-12.vrack.vsphere.local	10.0.0.111	bringup-networkpool	Active	Unassigned	-	1%	7%	Hybrid	-	2
<input type="checkbox"/>	esxi-13.vrack.vsphere.local	10.0.0.112	bringup-networkpool	Active	Unassigned	-	1%	7%	Hybrid	-	2

1 Hosts 1 - 4 of 4

- In the hosts table, select the host(s) you want to decommission.
- Click **Decommission Selected Hosts**.
- Activate or deactivate **Skip failed hosts during decommissioning**.

This setting is enabled by default. When enabled, hosts that fail are skipped and decommissioning continues on the remaining hosts.

- Click **Confirm**.

Re-image the decommissioned host before commissioning it again and adding it to a workload domain.

ESXi Lockdown Mode

You can activate or deactivate normal lockdown mode in VMware Cloud Foundation to increase the security of your ESXi hosts.

To activate or deactivate normal lockdown mode in VMware Cloud Foundation, you must perform operations through the vCenter Server. For information on how to activate or deactivate normal lockdown mode, see [vSphere Security](#).

You can activate normal lockdown mode on a host after the host is added to workload domain. VMware Cloud Foundation creates service accounts that can be used to access the hosts. Service accounts are added to the Exception Users list during the bring-up or host commissioning. You can rotate the passwords for the service accounts using the password management functionality in the SDDC Manager UI.

Managing vSphere Lifecycle Manager Images in VMware Cloud Foundation

A vSphere Lifecycle Manager image represents a desired software specification to be applied to all ESXi hosts in a vSphere cluster. When you set up a vSphere Lifecycle Manager image, you define the full software stack that you want to run on the ESXi hosts in a vSphere cluster: the ESXi version, additional VMware software, and vendor and third-party software such as firmware and drivers. Using a single image to manage all ESXi hosts in a vSphere cluster ensures cluster-wide host image homogeneity.

vSphere Lifecycle Manager images are made available by leveraging vSphere Lifecycle Manager, a vCenter Server service. This service is integrated with VMware Cloud Foundation and enables centralized and simplified lifecycle management of ESXi hosts. When you use vSphere Lifecycle Manager images, you follow one workflow and use the

same ESXi image for all software lifecycle related operations: install, upgrade, update, and patching, which significantly simplifies the life cycle management process.

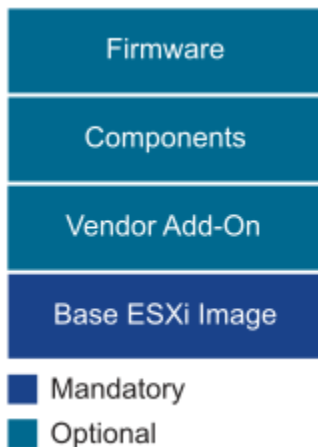
For more information on vSphere Lifecycle Manager images, see [Using vSphere Lifecycle Manager Images](#).

vSphere Lifecycle Manager Image Components

A vSphere Lifecycle Manager image can include four elements.

- **ESXi base image**
The base image contains an image of VMware ESXi Server and additional components, such as drivers and adapters that are necessary to boot a server. Base images have a user-readable name and a unique version.
- **Vendor Add-On**
A vendor add-on is a collection of software components that OEMs create and distribute for ESXi hosts. The vendor add-on can contain drivers, patches, and solutions.
- **Component**
A component is the smallest discrete unit in an image, created and published by third-party software vendors.
- **Firmware**
Firmware refers to firmware and drivers add-on, a special type of vendor add-on designed to assist in the firmware update process. It contains firmware for a specific server type and corresponding drivers. To add a firmware and drivers add-on to your image, you must install the hardware support manager plug-in provided by the hardware vendor for the hosts in the respective cluster.

The ESXi base image is required for setting up a vSphere Lifecycle Manager image. All other elements are optional.



vSphere Lifecycle Manager Images in VMware Cloud Foundation

vSphere Lifecycle Manager images must be created in vSphere and then imported to VMware Cloud Foundation. Unlike vSphere where images are managed per cluster, VMware Cloud Foundation allows you to manage all images in a single place and re-use them for clusters across workload domains.

You can create an image either on the management domain vCenter Server, or a vCenter Server external to VMware Cloud Foundation. Since an image is created for a cluster of hosts, you must create a cluster enabled for vSphere Lifecycle Manager images and then create an image for that cluster. While creating the image, you specify the ESXi version and can also select vendor add-ons, components, or firmware for the image.

If you created the image in a vCenter Server managed by VMware Cloud Foundation, you can extract the image from the vSphere cluster. If you created the image in an external vCenter Server, you export the image specification and component files from vSphere to your local computer. You then import these files to VMware Cloud Foundation.

After you import an image to VMware Cloud Foundation, it can be customized when it is applied to a cluster. For example, say your image included only the base ESXi image. When you use this image to create a VI workload domain, NSX components are added to the default cluster during the domain creation. You may also add vendor-adds, components, or firmware to the workload domain cluster. You can now extract the updated image from this cluster and reuse it for other similar clusters in this workload domain or in other workload domains.

vSphere Lifecycle Manager Image Workflow



vSphere Lifecycle Manager Options for Workload Domains

Prior to VMware Cloud Foundation 5.2.1, you have two vSphere Lifecycle Manager options for workload domains. You select an option while creating the workload domain. All clusters in the workload domain are managed using the selected option.

- vSphere Lifecycle Manager images automate the lifecycle management of your VMware Cloud Foundation environment including firmware by using cluster images.
- vSphere Lifecycle Manager baselines also automates the lifecycle management of your VMware Cloud Foundation environment, but firmware updates are manual.

VMware Cloud Foundation 5.2.1 and later support both vSphere Lifecycle Manager options within the same workload domain. For example, you could have one cluster within the workload domain that uses vSphere Lifecycle Manager images and another cluster in the same workload domain that uses vSphere Lifecycle Manager baselines.

Create a vSphere Lifecycle Manager Image

vSphere Lifecycle Manager images are created using the vSphere Client. You can create an image either on the management domain vCenter Server, or a vCenter Server external to VMware Cloud Foundation.

If you want to add firmware to the vSphere Lifecycle Manager image, you must install the Hardware Support Manager from your vendor. See [Firmware Updates](#).

You first create an empty cluster in vSphere Client and then configure a vSphere Lifecycle Manager image on that vSphere cluster. During the creation of an image, you define the ESXi version and can optionally add vendor add-ons, components, and firmware.

1. Log in to the vSphere Client of the vCenter Server where you will create the vSphere Lifecycle Manager Image.
2. Import the depot ZIP file for the version of ESXi into the vSphere Lifecycle Manager depot.





You need to perform this step only if the ESXi version listed in the Bill of Materials section of the *VMware Cloud Foundation Release Notes* is not present in the vSphere Lifecycle Manager depot.



1. Download the ESXi depot ZIP file that matches the ESXi version in the Bill of Materials section of the *VMware Cloud Foundation Release Notes*.
2. Click **Menu** > **Lifecycle Manager**.
3. On the Lifecycle Manager page, click **Actions** > **Import Updates**.
4. Click **Browse**, locate the ESXi ZIP file, and click **Open**.

3. Create an empty cluster for the vSphere Lifecycle Manager image. You do not need to add any hosts to this cluster.

1. In the vSphere Client, click **Menu > Hosts and Clusters**.
2. In the navigation pane, click **vCenter > Hosts and Clusters**.
3. Right-click the datacenter and click **New Cluster**.
4. Enter a name for the cluster.
As a best practice, name this cluster *ClusterForImage*. You can keep this cluster for creating additional vSphere Lifecycle Manager images for upgrades.
5. Select **Manage all hosts in this cluster with a single image** and choose **Compose a new image**.

Basics

Name	<u>ClusterForImage</u>
Location	 SDDC-Datacenter
 vSphere DRS	<input type="checkbox"/>
 vSphere HA	<input type="checkbox"/>
vSAN	<input type="checkbox"/> Enable vSAN ESA 

- Manage all hosts in the cluster with a single image 
- Choose how to set up the cluster's image
- Compose a new image
- Import image from an existing host in the vCenter inventory
- Import image from a new host
- Manage configuration at a cluster level 

CANCEL

NEXT

6. Click **Next**.
7. In the **Image Setup** section, select the ESXi software you uploaded in step 2.
8. In Vendor Addon, select the vendor add-on and click **Next**.
9. Review the image details and click **Finish**.

The new cluster and vSphere Lifecycle Manager image are created.

Do one of the following to make a vSphere Lifecycle Manager image available in VMware Cloud Foundation:

- If you created a vSphere Lifecycle Manager image in a vCenter Server managed by VMware Cloud Foundation, extract the image. See [Extract a vSphere Lifecycle Manager Image](#).
- If you created a vSphere Lifecycle Manager image in an external vCenter Server, export the image from vCenter Server and import it to VMware Cloud Foundation. See [Export a vSphere Lifecycle Manager Image](#).

Export a vSphere Lifecycle Manager Image

If you created a vSphere Lifecycle Manager image in an external vCenter Server, you export the image from vCenter Server to your local computer.

Create a vSphere Lifecycle Manager image in the vSphere Client. For more information, see [Create a vSphere Lifecycle Manager Image](#).

The following files need to be exported from vCenter Server.

Format	Content
JSON	Image specification
JSON	Cluster specification
ISO	ESXi image (optional) You can use the ISO file for imaging additional hosts that you bring into VMware Cloud Foundation.
ZIP	Image components

1. Log in to the vSphere Client of the vCenter Server where you created the vSphere Lifecycle Manager image.
2. Export the image specification JSON, ISO, and ZIP files from vSphere to your local computer.
 1. In the vSphere Client, click **Host and Clusters**.
 2. In the navigation pane, expand **vCenter > Datacenter** and select the vSphere Lifecycle Manager image you want to export.
 3. On the Updates tab, click **Hosts > Image**.
 4. Click the horizontal ellipsis icon and select **Export**.
 5. Click **JSON** and then click **Export**.
This downloads the image specification JSON file.
 6. Repeat steps c and d for the ISO and ZIP file formats.
You can use the ISO file for imaging additional hosts that you bring into VMware Cloud Foundation. Exporting this file is optional.

NOTE

Do not rename these files.

3. Retrieve the vSphere cluster ID using vSphere APIs.
 - a) Click **Menu > Developer Center**.
 - b) Click the **API Explorer** tab.
 - c) Ensure that **Select Endpoint** has the local vCenter Server selected and **Select API** is set to vcenter.
 - d) Expand the **cluster** section.
 - e) Expand **GET /api/vcenter/cluster** and click **Execute**.
 - f) Under Response, expand **VcenterClusterSummary (...)** for the cluster where you created the image and copy the `cluster` value.
4. Retrieve the cluster JSON specification and download to the local computer using vSphere APIs.
 - a) In the **API Explorer** tab of the vCenter Server, ensure that **Select Endpoint** has the local vCenter Server selected and **Select API** is set to esx.
 - b) Expand the **settings/clusters/software** section.
 - c) Expand **GET /api/esx/settings/clusters/{cluster}/software**.
 - d) In the cluster parameter area, paste the `cluster` value you had retrieved in step 3 and click **Execute**.

- e) Under Response next to EsxSettingsSoftwareInfo, click the download arrow to download the JSON specification.

Creating a vSphere Lifecycle Manager Image in VMware Cloud Foundation

A vSphere Lifecycle Manager image must be made available in VMware Cloud Foundation before it can be applied to a cluster when creating a VI workload domain or adding a vSphere cluster. You can either extract the vSphere Lifecycle Manager image from an existing vSphere Lifecycle Manager enabled workload domain cluster or import a vSphere Lifecycle Manager image that has been created and exported from an external vCenter Server.

A vSphere Lifecycle Manager image can be customized when it is applied to a vSphere cluster in VMware Cloud Foundation. For example, say your vSphere Lifecycle Manager image included only the base ESXi image. When you use this vSphere Lifecycle Manager image to create a VI workload domain, NSX components are added to the default cluster during the VI workload domain creation. You may also add vendor-adds, components, or firmware to the VI workload domain cluster. You can now extract the updated image from this cluster and reuse it for other similar clusters in this VI workload domain or in other VI workload domains.

Extract a vSphere Lifecycle Manager Image

You can extract a vSphere Lifecycle Manager image from a vCenter Server managed by VMware Cloud Foundation.

A vSphere Lifecycle Manager image must have been created in vSphere. For more information, see [Create a vSphere Lifecycle Manager Image](#).

1. From the navigation bar, click **Lifecycle Management > Image Management**.
2. Click the **Import Image** tab.
3. In the Option 1 section, select a workload domain.

The screenshot shows a user interface for selecting a cluster image. At the top, there is a tab labeled 'Option 1' and a title 'Extract a Cluster Image'. Below the title is a descriptive text: 'Extract a cluster image assigned to a cluster that was updated in vCenter.' There are two dropdown menus: 'Select Workload Domain' and 'Select Cluster'. At the bottom of the selection area is a large button with a refresh icon and the text 'EXTRACT CLUSTER IMAGE'.

4. Select the cluster from which you want to extract the vSphere Lifecycle Manager image.
5. Click **Extract Cluster Image**.

The extracted cluster image is displayed in the Available Images tab and can be used for a new VI workload domain or a new cluster in a VI workload domain enabled for vSphere Lifecycle Manager images.

Import a vSphere Lifecycle Manager Image

If you had exported a vSphere Lifecycle Manager image from an external vCenter Server to your local computer, you can import it to VMware Cloud Foundation.

Verify the exported files for a vSphere Lifecycle Manager image are available to be imported. See [Export a vSphere Lifecycle Manager Image](#).

The following files are imported to VMware Cloud Foundation.

Format	Content
JSON	Image specification
JSON	Cluster specification
ISO	ESXi image (optional) In addition to the ESXi base image, the ISO file includes vendor-adds, components, and firmware associated with the vSphere Lifecycle Manager image. You can use this ISO file for imaging additional hosts that you bring into VMware Cloud Foundation.
ZIP	Image components

1. In SDDC Manager, click **Lifecycle Management > Image Management**.
2. Under Option 2 **Import a Cluster Image** section, select the JSON, ZIP, and ISO files from your local computer. The ISO file is optional.

Option 2 Import a Cluster Image

Import the image files for an exported cluster image from an external vCenter.

SELECT FILE Select Cluster Settings JSON File

SELECT FILE Select Software Spec JSON File

SELECT FILE Select ZIP File

SELECT FILE Select ISO File (Optional)

Cluster Image Name

↑ UPLOAD IMAGE COMPONENTS

A vSphere Lifecycle Manager image is referred to as a cluster image on the SDDC Manager UI. Ensure that you upload the image specification JSON and cluster specification JSON files in the correct fields. Mixing these files up will result in a validation error.

3. Enter a name for the vSphere Lifecycle Manager image and click **Upload Image Components**.

The imported vSphere Lifecycle Manager image is displayed in the Available Images tab and can be used for a new VI workload domain or a new vSphere cluster in a VI workload domain that is enabled for vSphere Lifecycle Manager images.

Firmware Updates

You can use vSphere Lifecycle Manager images to perform firmware updates on the ESXi hosts in a cluster. Using a vSphere Lifecycle Manager image simplifies the host update operation. With a single operation, you update both the software and the firmware on the host.

To apply firmware updates to hosts in a cluster, you must deploy and configure a vendor provided software module called hardware support manager. The deployment method and the management of a hardware support manager is determined by the respective OEM. For example, the hardware support manager that Dell EMC provides is part of their host management solution, OpenManage Integration for VMware vCenter (OMIVV), which you deploy as an appliance. See [Deploying Hardware Support Managers](#).

You must deploy the hardware support manager appliance on a host with sufficient disk space. After you deploy the appliance, you must power on the appliance virtual machine, log in to the appliance as an administrator, and register the appliance as a vCenter Server extension. Each hardware support manager has its own mechanism of managing firmware packages and making firmware add-ons available for you to choose. For detailed information about deploying, configuring, and managing hardware support managers, refer to the vendor-provided documentation.

View vSphere Lifecycle Manager Images

You can view the cluster images available in VMware Cloud Foundation.

1. From the navigation bar, click **Lifecycle Management > Image Management**.
2. Click the **Available Images** tab.

The screenshot shows the 'Available Images' tab in the vSphere Lifecycle Manager interface. It features a search bar for 'Search Image Name' and filters for 'Filter By Version' (set to 'All ESXi Versions') and 'Filter By Vendor' (set to 'All Vendors'). A table lists the available images with columns for image name, release date, size, and creator.

Image Name	Released	Size	Created by
8.0 U1a - 21813344	Jun 8, 2023	0 bytes	

Below the table, there are two columns: 'ESXi Version' with a value of '8.0.1-21813344' and 'Components' with a value of 'No component'. Both columns have an information icon (i).

All vSphere Lifecycle Manager images available in your VMware Cloud Foundation environment are displayed. You can search by image name, or filter by ESXi version or vendor.

Managing Storage in VMware Cloud Foundation

To create and manage a workload domain, VMware Cloud Foundation requires at least one shared storage type for all ESXi hosts within a cluster. This initial shared storage type, known as principal storage, is selected during the creation of a workload domain or cluster in SDDC Manager. Additional shared storage, known as supplemental storage, can be added using the vSphere Client after a cluster has been created.

Principal Storage

When you create a VI workload domain or add a cluster to a workload domain in SDDC Manager, you must select the initial shared storage type. This initial shared storage type is known as principal storage. Once created, the principal storage type for a cluster cannot be changed. However, a VI workload domain can include multiple clusters with unique principal storage types.

VMware Cloud Foundation supports the following types of principal storage:

- vSAN
 - vSAN Original Storage Architecture (vSAN OSA)
 - vSAN Express Storage Architecture (vSAN ESA)
 - vSAN Max (requires vSAN ESA HCL compatible hosts)

NOTE

You cannot convert vSAN OSA to vSAN ESA or vice versa.

- Network File System (NFS)
- VMFS on FC (Fibre Channel)
- VMware vSphere Virtual Volumes (vVols)

NOTE

vVols supports FC, NFS, and iSCSI storage protocol types.

Supplemental Storage

Additional shared storage, known as supplemental storage, can be manually added or removed using the vSphere Client after a cluster has been created. All supplemental storage must be listed in the VMware Compatibility Guide. Multiple supplemental storage types can be presented to a cluster in the management domain or any VI workload domain.

VMware Cloud Foundation supports using the vSphere Client to add the following datastore types to a cluster:

- vSphere VMFS
- NFS
- VMware vSphere Virtual Volumes (vVols)

NOTE

vVols supports FC, NFS, and iSCSI storage protocol types.

vSAN Storage with VMware Cloud Foundation

vSAN is the preferred principal storage type for VMware Cloud Foundation. It is an enterprise-class storage integrated with vSphere and managed by a single platform. vSAN is optimized for flash storage and can non-disruptively expand capacity and performance by adding hosts to a cluster (scale-out) or by adding disks to a host (scale-up).

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	Yes	Yes	Yes
Supplemental	No	No	No

Prerequisites for vSAN Storage

To create a VI workload domain that uses vSAN as principal storage you must ensure the following:

- A minimum of three ESXi hosts that meet the vSAN hardware, cluster, software, networking and license requirements. For information, see the [vSAN Planning and Deployment Guide](#).
- The hosts must be in the SDDC Manager inventory. See [Commission Hosts](#).
- A network pool that includes details for the vMotion and vSAN networks that will be used for the cluster. See [Network Pool Management](#).
- A valid vSAN license. See [Managing License Keys in](#) . You cannot use vSAN ESA or vSAN Max without a qualifying license.

In some instances SDDC Manager may be unable to automatically mark the host disks as capacity. Follow the Mark Flash Devices as Capacity Using ESXCLI procedure in the [vSAN Planning and Deployment Guide](#).

Procedures for vSAN Storage

- To use vSAN as principal storage for a new VI workload domain, see [Deploy a VI Workload Domain Using the SDDC Manager UI](#).
- To use vSAN as principal storage for a new cluster, see [Add a vSphere Cluster to a Workload Domain Using the SDDC Manager UI](#).

vSAN Original Storage Architecture (OSA)

With vSAN OSA, each host that contributes storage devices to the vSAN datastore must provide at least one device for flash cache and at least one device for capacity. The devices on the contributing host form one or more disk groups. Each disk group contains one flash cache device, and one or multiple capacity devices for persistent storage. Each host can be configured to use multiple disk groups.

vSAN OSA clusters can mount a remote datastore from other vSAN OSA clusters.

vSAN Express Storage Architecture (ESA)

With vSAN ESA, all storage devices claimed by vSAN contribute to capacity and performance. Each host's storage devices claimed by vSAN form a storage pool. The storage pool represents the amount of caching and capacity provided by the host to the vSAN datastore.

vSAN ESA clusters can mount a remote datastore from vSAN Max clusters or other vSAN ESA clusters.

To use vSAN ESA, you need:

- A direct or proxy internet connection OR a downloaded copy of the vSAN HCL JSON file

NOTE

SDDC Manager will keep the HCL file updated if it has direct or proxy internet connection.

- ESXi host disks to support vSAN ESA
- A vLCM image to manage clusters.
 - If you need to create a new cluster image, go to the management domain's vCenter to create the image.
 - If there are no vLCM images available, go to Image Management (**Lifecycle Management > Image Management > Import Image**) to extract or import a cluster image. See [Extract a vSphere Lifecycle Manager Image](#) and [Import a vSphere Lifecycle Manager Image](#) for more information.
 - See the [vSAN ESA VCG](#) for information about compatible hardware.

vSAN Max

vSAN Max is a fully distributed, scalable, shared storage solution. Storage resources are disaggregated from compute resources, so you can scale storage and compute resources independently. A vSAN Max cluster acts as a server cluster that only provides storage. You can mount its datastore to vSAN compute clusters or vSAN ESA clusters.

vSAN Max uses vSAN Express Storage Architecture and high-density vSAN Ready Nodes for increased capacity and performance.

Since vSAN Max uses vSAN ESA, it has the same requirements listed above.

vSAN Compute Clusters

A vSAN compute cluster is a vSphere cluster with a small vSAN element that enables it to mount a remote datastore. The hosts in a compute cluster do not have local storage. A compute cluster can only mount a remote datastore from a cluster within the same workload domain.

A vSAN compute cluster can mount a datastore from one of the following cluster types:

- vSAN OSA
- vSAN ESA
- vSAN Max

Once you mount a remote datastore on a vSAN compute cluster, you can only mount additional datastores of the same cluster type.

NOTE: Datastores on clusters created outside of VMware Cloud Foundation cannot be mounted on VCF-created clusters. Likewise, clusters created outside of VMware Cloud Foundation cannot mount a datastore from a VCF-created cluster.

NFS Storage with VMware Cloud Foundation

An NFS client built into ESXi uses the Network File System (NFS) protocol over TCP/IP to access a designated NFS volume that is located on a NAS server. NFS Storage with VMware Cloud Foundation can mount a single NFS volume to each ESXi host where it can be used to store virtual disks or as a central repository for ISO images and virtual machine templates.

NFS is generally used as a supplemental storage option but can be used as principal storage with VI workload domains. NFS can also be used as a principal storage in a management domain converted from vSphere infrastructure.

VMware Cloud Foundation only supports NFS protocol version 3 when used as principal storage. Supplemental storage can use either vSphere supported NFS protocol version 3 or 4.1. Although NFS 3 and NFS 4.1 can coexist on the same host, you cannot use different NFS versions to mount the same datastore on different hosts.

Only a single NFS volume is mounted during the creation of a new workload domain or cluster, additional NFS volumes can be mounted as supplemental storage after the VI workload domain or cluster has been created.

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	No	Only for a management domain converted from vSphere infrastructure	Yes
Supplemental	Yes	Yes	Yes

An additional VMkernel port is created during the VI workload domain or cluster creation for NFS. If the NFS server is on the same subnet as the NFS VMkernel port this will be used for NFS traffic. If the NFS server is on a different subnet the NFS traffic is routed over the ESXi management VMkernel port.

Prerequisites for NFS Storage

- A minimum of three ESXi hosts marked with the NFS storage must be in the SDDC Manager inventory. See [Commission Hosts](#).

NOTE

If you are using NFS as principal storage, and your VI workload domain is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- An NFS Server that meets the NFS Storage Guidelines and Requirements found in the [vSphere Storage Guide](#).
- A Network Pool that includes details for the vMotion and NFS networks that will be used for the cluster. See [Network Pool Management](#).
- The FQDN and the mount point folder name of the NFS server.
- For L3-routed NFS the ESXi Management network must be able to reach the NFS server network and have access to the NFS server.

Procedures for NFS Storage

- To use NFS as principal storage for a new VI workload domain, see [Deploy a VI Workload Domain Using the SDDC Manager UI](#).
- To use NFS as principal storage for a new cluster, see [Add a vSphere Cluster to a Workload Domain Using the SDDC Manager UI](#).

- To use NFS as supplemental storage follow the *Create an NFS Datastore* procedure in the [vSphere Storage Guide](#).

Fibre Channel Storage with VMware Cloud Foundation

Fibre Channel (FC) is a storage protocol that the SAN uses to transfer data traffic from ESXi hosts to shared storage. The protocol packages SCSI commands into FC frames. To connect to the FC SAN, the ESXi host uses Fibre Channel host bus adapters (HBAs).

Fibre Channel can be used as supplemental storage for the management domain and consolidated workload domains, however it can be used as principal storage for VI workload domains and can also be used a principal storage in a management domain converted from vSphere infrastructure.

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	No	Only for a management domain converted from vSphere infrastructure	Yes
Supplemental	Yes	Yes	Yes

Prerequisites for FC Storage

- A minimum of three ESXi hosts. Review the ESXi Fibre Channel SAN Requirements in the [vSphere Storage Guide](#).

NOTE

If you are using VMFS on FC as principal storage, and your VI workload domain is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- A pre-created VMFS datastore.
- The hosts must be in the SDDC Manager inventory. See [Commission Hosts](#).

IMPORTANT

The VMFS datastore must be mounted on all the hosts before commissioning.

- A network pool that includes details for the vMotion network that will be used for the cluster. See [Network Pool Management](#).

Procedures for FC Storage

- To use Fibre Channel as principal storage for a new VI workload domain, see [Deploy a VI Workload Domain Using the SDDC Manager UI](#).
- To use Fibre Channel as principal storage for a new cluster, see [Add a vSphere Cluster to a Workload Domain Using the SDDC Manager UI](#).
- To use Fibre Channel as supplemental storage, see the [vSphere Storage Guide](#).

HCI Mesh with VMware Cloud Foundation

HCI Mesh is a software-based approach for disaggregation of compute and storage resources in vSAN. HCI Mesh brings together multiple independent vSAN clusters by enabling cross-cluster use of remote datastore capacity within vCenter Server. HCI Mesh enables you to efficiently use and consume data center resources, which provides simple storage management at scale.

VMware Cloud Foundation 4.2 and later supports sharing remote datastores with HCI Mesh for VI workload domains.

You can create HCI Mesh by mounting remote vSAN datastores on vSAN clusters and enable data sharing from the vCenter Server or the SDDC Manager. It can take up to 5 minutes for the mounted remote vSAN datastores to appear in the SDDC Manager UI.

NOTE

HCI Mesh does not support remote vSAN datastores on stretched clusters.

For more information on sharing remote datastores with HCI Mesh, see [Sharing Remote Datastores with HCI Mesh](#).

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	No	No	Yes
Supplemental	Yes	Yes	Yes

HCI Mesh Compute-only Clusters

VMware Cloud Foundation now supports HCI Mesh as a principal storage type with compute-only clusters. With this feature, the hosts in an HCI Mesh compute-only client cluster do not need local storage. They can mount remote datastores from a vSAN cluster located within the same data center.

VMware Cloud Foundation can only support the use of cross-cluster HCI Mesh within a single Workload Domain. You cannot create a cluster in Workload Domain 1 and share or remotely connect to a vSAN datastore that exists in a cluster in Workload Domain 2.

Also, this feature only supports VMware Cloud Foundation-created clusters. For example, vSphere clusters created outside of VMware Cloud Foundation cannot be used to provide storage or remotely connect to an existing VMware Cloud Foundation-created cluster that is providing storage through HCI mesh.

vVols Storage with VMware Cloud Foundation

Virtual Volumes (vVols) is an integration and management framework that virtualizes SAN/NAS arrays. vVols shares a common storage operational model with vSAN and uses storage policy-based management (SPBM). Descriptive policies are assigned at the VM or VMDK level and are applied or changed within minutes. (vVols) is a storage type for VI workload domains and clusters in VMware Cloud Foundation.

With vVols, an individual virtual machine disk (VMDK) becomes a unit of storage management. The storage hardware gains control over the virtual disk content, layout, and its management.

vVols are exported to ESXi host(s) through a small set of Protocol Endpoints (PE). The storage system provides protocol endpoints that are discoverable on the physical storage fabric. Protocol Endpoints are part of the physical storage fabric, and they establish a data path from virtual machines to their respective vVols on demand. ESXi hosts use Protocol Endpoints to connect to virtual volumes on the storage system. Protocol Endpoints provide uniform access to both SAN (FC, iSCSI) and NAS (NFS) storage, regardless of the storage protocol selected.

vSphere APIs for Storage Awareness (VASA) allow the storage system to become aware of vVols and their associations with the relevant virtual machines. Through VASA, vSphere and the underlying storage system establishes a two-way out-of-band communication to perform data services and offload certain virtual machine operations to the storage system. For example, operations such as snapshots and clones can be offloaded.

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	No	No	Yes
Supplemental	Yes	Yes	Yes

NOTE

vVols supports FC, NFS, and iSCSI storage protocol types.

Prerequisites for vVols Storage

Before using vVols as principal or supplemental storage, review the *Working with VMware vSphere Virtual Volumes (vVols)* section in the [vSphere Storage Guide](#)

- Prepare the storage system for vVols
 - The storage system or storage array that you use must support vVols and integrate with the vSphere components through vSphere APIs for Storage Awareness (VASA). The storage array must support thin provisioning and snapshots.
 - The vVols storage provider must be deployed.
 - The following components must be configured on the storage side:
 - Protocol endpoints.
 - Storage containers.
 - Storage profiles.
 - Replication configurations if you plan to use vVols with replication.
- Prepare the ESXi Hosts for vVols
 - A minimum of three hosts marked with the vVols must be in the SDDC Manager inventory. See [Commission Hosts](#).

NOTE

If you are using vVols as principal storage, and your VI workload domain is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- For end-to-end vVols support, HBA drivers need to support vVols-based devices. Review the I/O Devices section of the VMware Compatibility Guide, to verify Secondary LUNID support which enables vVols.
- Make sure to follow appropriate setup guidelines for the type of storage you use, Fibre Channel, iSCSI, or NFS. If necessary, install and configure storage adapters on your ESXi hosts.
 - If you use Fibre Channel, all ESXi hosts must have access to the Fibre Channel Array. Work with your storage administrator to ensure that Fibre Channel Switch Zoning has been properly configured, and that the ESXi hosts have been registered with the Storage Array.
 - If you use iSCSI, activate the software iSCSI adapters on your ESXi hosts. Configure Dynamic Discovery and enter the IP address of your vVols storage system.
 - If you use NFS, all ESXi hosts must be able to reach the NFS Server from the NFS Network assigned in the IP Pool.
- A Network Pool that includes details for the vMotion network that will be used for the cluster. See [Network Pool Management](#).
- Add a VASA Provider.

Procedures for vVols Storage

- To use vVols as principal storage for a new VI workload domain, see [Deploy a VI Workload Domain Using the SDDC Manager UI](#).
- To use vVols as principal storage for a new cluster, see [Add a vSphere Cluster to a Workload Domain Using the SDDC Manager UI](#).
- To use vVols as supplemental storage follow the "Configure Virtual Volumes" procedure in the [vSphere Storage Guide](#).

Add a VASA Provider

To use vVols storage in VMware Cloud Foundation, add a VASA provider.

Before you add the VASA provider details in VMware Cloud Foundation, ensure that you have set up the VASA provider. For information about setting up a VASA provider, see Before You Enable vVols topic in vSphere Storage.

1. In the navigation pane, click **Administration** › **Storage Settings**.
2. Click **Add VASA Provider**.

3. Enter a name for the VASA provider.
4. Enter a URL for the VASA provider.
5. Enter VASA provider user credentials.
6. Enter VASA provider storage container details.
7. Click **Save**.

NOTE

You can add additional VASA user credentials and storage containers.

View a VASA Provider

You can view the VASA provider details in SDDC Manager.

1. In the navigation pane, click **Administration** › **Storage Settings**.
2. In the VASA provider table, click the name of the VASA provider.

Edit a VASA Provider

You can edit the existing VASA provider details.

1. In the navigation pane, click **Administration** › **Storage Settings**.
2. Click the vertical ellipsis (three dots) in the VASA provider row you want to edit and then click **Edit**.
3. Edit the VASA provider details.

Delete a VASA Provider

You can delete a VASA provider through the SDDC Manager.

1. In the navigation pane, click **Administration** › **Storage Settings**.
2. Click the vertical ellipsis (three dots) in the VASA provider row you want to edit and then click **Delete**.
3. Confirm your choice by once again clicking **Delete**.

Converting or Importing Existing vSphere Environments into VMware Cloud Foundation

VMware Cloud Foundation 5.2 introduces a new CLI tool, the VCF Import Tool, to convert or import an existing vSphere environment that is not currently managed by VCF into VMware Cloud Foundation.

If you do not already have SDDC Manager deployed, you can deploy it on an existing vSphere environment and use the VCF Import Tool to convert that environment to the VMware Cloud Foundation management domain.

If SDDC Manager is already deployed, you can use the VCF Import Tool to import existing vSphere environments as VI workload domains.

In addition to importing and converting vSphere environments, you can use the VCF Import Tool to deploy VLAN-backed NSX and sync the SDDC Manager inventory.

- [Deploy NSX Manager for Workload Domains](#)
- [Manage Workload Domain Configuration Drift Between vCenter Server and SDDC Manager](#)

Supported Scenarios for Converting or Importing vSphere Environments to VMware Cloud Foundation

There are two main scenarios for using the VCF Import Tool, depending on whether or not you already have SDDC Manager deployed.

Supported Scenarios

Scenarios	Description
Scenario A: Convert	Your environment does not have SDDC Manager and you want to convert existing vSphere infrastructure to the VMware Cloud Foundation management domain.
Scenario B: Import	Your environment does have SDDC Manager deployed and you want to import existing vSphere infrastructure as VI workload domains.

Glossary of Terms for Converting/Importing vSphere Environments

Term	Description
BOM	Bill of materials. List of supported software versions for a given VMware Cloud Foundation release
Co-located	vCenter Server virtual machine is located on a cluster that it manages
Default cluster	The first cluster in the management workload domain where the management vCenter virtual machine is hosted.
ELM	vCenter Single Sign-On Enhanced Linked Mode (Multiple vCenter Servers in the same SSO domain)
SSO	vCenter Single Sign-On

Considerations Before Converting or Importing Existing vSphere Environments into VMware Cloud Foundation

This section describes the supported and unsupported configurations for converting an existing vSphere environment into a VMware Cloud Foundation management domain, or importing an existing vSphere environment as a VI workload domain.

The requirements and supported configurations vary slightly across different versions. Confirm the versions of the VCF Import Tool and SDDC Manager that you are running support your use case.

Considerations by Domain Type

Category	Management Domain Considerations	VI Workload Domain Considerations	Recommendation
VMware Cloud Foundation Software BOM Alignment	VCF Import Tool 5.2 requires: <ul style="list-style-type: none"> vCenter Server 8.0 Update 3a or later ESXi 8.0 Update 3 or later VCF Import Tool 5.2.1 and 5.2.1.2 require: <ul style="list-style-type: none"> vCenter Server 8.0 Update 3c or later 	All versions of the VCF Import Tool support: <ul style="list-style-type: none"> vCenter Server 7.0 Update 3h or later ESXi 7.0 Update 3g or later 	Upgrade the domain to at least the minimum required versions.

Table continued on next page

Continued from previous page

Category	Management Domain Considerations	VI Workload Domain Considerations	Recommendation
	<ul style="list-style-type: none"> ESXi 8.0 Update 3b or later 		<p>CAUTION</p> <p>vCenter Servers originally deployed at a version below 6.5 may encounter an issue upgrading to 8.0U3. See KB370882 for more details.</p>
Ports & Protocols	<p>Must align with https://ports.esp.vmware.com/home/VMware-Cloud-Foundation</p> <p>CAUTION</p> <p>vCenter Server must be using port 443.</p>		<p>Custom ports for vCenter Server are not currently supported for import. Please wait for a future version of VCF that will support importing vCenter Server using custom ports.</p>
vCenter Server VM Location	Must be co-located	Must be located in the management domain, or co-located	<p>Move the vCenter Server VM to a supported location. See this for more information on cross vCenter vMotion.</p>
Single Sign-On	SSO domain names for imported environments do not need to be unique within a VMware Cloud Foundation instance		
	Each SSO domain should contain only a single vCenter Server. ELM is not supported		<p>Remove the vCenter instance from the ELM ring, creating multiple SSO domains</p>
Cluster - Storage	<ul style="list-style-type: none"> Default cluster must be one of vSAN, NFS v3, VMFS-FC, or VMFS-FCoE (only supported with VCF 5.2.1 and later). NFS 4.1, VVOLs, and native iSCSI are not supported 		<p>Select a cluster with primary storage from the supported list</p>
	Clusters cannot be stretched vSAN		<p>vSAN stretched clusters are not currently supported. Please wait for a future version of VCF that will support importing vSAN stretched clusters.</p>
	<ul style="list-style-type: none"> VCF 5.2: All clusters (vSAN, NFS v3, FC) must be 4 nodes minimum. VCF 5.2.1.x: 	<ul style="list-style-type: none"> When using vSAN, all clusters must be 3 nodes minimum. 	<p>Expand the cluster to the minimum number of nodes for the relevant storage type</p>

Table continued on next page

Continued from previous page

Category	Management Domain Considerations	VI Workload Domain Considerations	Recommendation
	<ul style="list-style-type: none"> When using vSAN, vLCM image clusters must be 3 nodes minimum and vLCM baseline clusters must be 2 nodes minimum. When using NFS or FC and vLCM images, the default cluster must be 2 nodes minimum. When using NFS or FC and vLCM baselines, the default cluster must be 4 nodes minimum. 	<ul style="list-style-type: none"> When using NFS or FC all clusters must be 2 nodes minimum 	
	When using vSAN, compression only (applicable for OSA) is not supported		Dedupe and compression are supported together. Either enable dedupe, or disable compression.
Cluster - Network	vCenter Server must not have an existing NSX instance registered		vCenter Servers with existing NSX registrations are not currently supported for import
	LACP support: <ul style="list-style-type: none"> VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1: Supported Earlier versions: Not supported 		For earlier versions of the VCF Import Tool and SDDC Manager, you can use teaming options available with vSphere Distributed Switch and N-VDS to provide load balancing and failover instead of LACP.
	Use vSphere Distributed Switches only. Standard or Cisco virtual switches are not supported. NOTE When using the VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1, your hosts can have a standard switch, as long they also have a vSphere distributed switch.		To move to a vSphere Distributed Switch. See this procedure .
	VMkernel IP addresses must be statically assigned		Move to statically assigned IP addresses
	Multiple VMkernels for a single traffic type (vSAN , vMotion) are not supported		Reconfigure to a single VMkernel per traffic type
	ESXi physical uplinks: <ul style="list-style-type: none"> VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1: ESXi hosts can have a different number of physical uplinks (minimum 2) assigned to a vSphere distributed switch. Each uplink must be a minimum of 10Gb. 		<ul style="list-style-type: none"> Use the VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1. Or, reconfigure uplinks accordingly.

Table continued on next page

Continued from previous page

Category	Management Domain Considerations	VI Workload Domain Considerations	Recommendation
	<ul style="list-style-type: none"> Earlier versions: ESXi hosts must have the same number of physical uplinks (minimum 2) assigned to a vSphere distributed switch. Each uplink must be a minimum of 10Gb. 		
	vSphere distributed switch teaming policies must match VMware Cloud Foundation standards		See here for VMware Cloud Foundation teaming policies
	Dedicated vMotion network must be configured		Configure a dedicated vMotion network
	vSphere distributed switch: <ul style="list-style-type: none"> VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1: Supports clusters that share a vSphere distributed switch. Earlier versions: Each cluster must have a dedicated vSphere distributed switch. 		<ul style="list-style-type: none"> Use the VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1 Or, reconfigure uplinks accordingly
Cluster - Compute	Clusters must not be VxRail managed		VxRail is not currently supported. Please wait for a future version of VCF that will support importing VxRail
	Clusters that use vSphere Configuration Profiles are not supported		vSphere Configuration Profiles were introduced as part of vSphere 8.0 and do not yet support NSX.
	All clusters must be running vSphere 8.0U3 or later	ESXi build number must be consistent within a cluster	Upgrade hosts to align build numbers
	DRS must be fully automated		Enable DRS and set to fully automated
	Standalone hosts: <ul style="list-style-type: none"> VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1: Supports standalone hosts, as long as the vCenter has a cluster that meets the import/convert requirements. <p style="text-align: center;">IMPORTANT You cannot perform any day-N operations on standalone hosts using SDDC Manager.</p> <ul style="list-style-type: none"> Earlier versions: Not supported 		Remove the standalone host from the vCenter Server inventory.
	Single host clusters: <ul style="list-style-type: none"> VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1: Supports single host clusters, as long as the vCenter has a cluster that meets the import/convert requirements. <p style="text-align: center;">IMPORTANT You cannot perform any day-N operations on standalone hosts using SDDC Manager.</p> <ul style="list-style-type: none"> Earlier versions: Not supported 		Use the VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1.

Table continued on next page

Continued from previous page

Category	Management Domain Considerations	VI Workload Domain Considerations	Recommendation
	vSphere Lifecycle Manager method <ul style="list-style-type: none"> VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1: Supports importing/converting a vCenter with a mix of vSphere Lifecycle Manager image and vSphere Lifecycle Manager baseline clusters. Earlier versions: Clusters in the same vCenter must all be managed using the same vSphere Lifecycle Manager method. You cannot import or convert a vCenter that includes clusters with a mix of vSphere Lifecycle Manager images and vSphere Lifecycle Manager baselines. 		Use the VCF Import Tool 5.2.1.2 with SDDC Manager 5.2.1.1.

Configuration Support by Domain Type

Table 181: Configuration Support by Domain Type

Configuration	Converted Management Domain	Imported VI Workload Domain	VCF Deployed Management Domain	VCF Deployed VI Workload Domain	Additional Information
Deployment with NSX VLAN-backed	Yes	Yes	No	No	To manually enable virtual networking for a converted/ imported workload domain see VMware Cloud Foundation: Enabling Virtual Networking on Imported Workload Domains .
NSX Edge deployment	Yes*	Yes*	Yes	Yes	* Requires virtual networking. See VMware Cloud Foundation: Enabling Virtual Networking on Imported Workload Domains .
AVN	Yes*	N/A	Yes	N/A	* Requires NSX edge cluster, which requires virtual networking. See VMware Cloud Foundation: Enabling Virtual

Table continued on next page

Continued from previous page

Configuration	Converted Management Domain	Imported VI Workload Domain	VCF Deployed Management Domain	VCF Deployed VI Workload Domain	Additional Information
					Networking on Imported Workload Domains.
Aria Lifecycle deployment in VCF aware mode	Yes*	N/A	Yes	N/A	* Requires virtual networking. See VMware Cloud Foundation: Enabling Virtual Networking on Imported Workload Domains.
Fibre channel storage as primary storage	Yes	Yes	No	Yes	
NFS storage as primary storage	Yes	Yes	No	Yes	
vCenter Appliances co-located in a workload domain cluster	Yes	Yes	Yes	No	
Non standard VCF networking, such as link aggregation to the hosts (LACP)	Yes*	Yes*	No	No	* Requires SDDC Manager 5.2.1.1 and VCF Import Tool 5.2.1.2.
Enhanced Link Mode (ELM)	No	No	Yes	Yes	
AVI Load Balancer	Yes*	Yes*	Yes	Yes	* Requires virtual networking. See VMware Cloud Foundation: Enabling Virtual Networking on Imported Workload Domains.
vSAN Stretched Cluster	No	No	Yes	Yes	
WCP Enabled clusters	No	No	Yes	Yes	
L3 vSphere Cluster create/ Add Host	No	No	No	Yes	

Table continued on next page

Continued from previous page

Configuration	Converted Management Domain	Imported VI Workload Domain	VCF Deployed Management Domain	VCF Deployed VI Workload Domain	Additional Information
vVOL enabled storage	No	No (New clusters can be added post import that utilize vVOL enabled storage)	No (Supplemental only)	Yes	
Multi-Region/DR	No	No	Yes	Yes	
VMware Validated Solutions	No	No	Yes	Yes	
ESXi Host Password management	No	No	Yes	Yes	
Add/remove host to a cluster	Yes*	Yes*	Yes	Yes	* Must be done in vCenter and then perform a sync operation

Network and Compute Requirements for Converting or Importing Existing vSphere Environments

This section describes the network requirements for converting or importing existing environments to VMware Cloud Foundation.

Network and Compute Requirements

Component	Network Requirement	Compute Requirement
SDDC Manager (for convert only)	<ul style="list-style-type: none"> 1 management IP address with corresponding DNS entry Must be routable to all components 	<ul style="list-style-type: none"> 4 vCPU 16Gb RAM 908Gb disk (Thin Provisioned)
NSX Manager	<ul style="list-style-type: none"> 4 IP addresses in the same subnet as the associated vCenter Server with corresponding DNS entries <ul style="list-style-type: none"> 3 x NSX Manager nodes 1 x NSX Manager VIP 	<ul style="list-style-type: none"> Compute requirements for NSX will vary depending on the appliance size chosen. Review the NSX documentation for sizing guidance. For management domain NSX Managers will be deployed in the management domain vCenter Server For VI workload domains NSX Managers will be deployed in the same vCenter as its corresponding vCenter Server VM NSX deployment requires a minimum of 3 hosts

Download Software for Converting or Importing Existing vSphere Environments

Before you can import or convert an existing vSphere environment, you must download the software.

Download the following software from the [Broadcom Support portal](#).

Table 182: Software Requirements for VMware Cloud Foundation 5.2

Software	Build Number	Details
VCF SDDC Manager Appliance (VCF-SDDC-Manager-Appliance-5.2.0.0-24108943.o va)	24108943	Only required for convert workflows.
VCF Import Tool (vcf-brownfield- import-5.2.0.0-24108578.tar. gz)	24108578	Required for convert, import, and sync workflows.
VMware Software Install Bundle - NSX_T_MANAGER 4.2.0.0 (bundle-124941.zip)	24105817	Required for convert or import workflows.

Table 183: Software Requirements for VMware Cloud Foundation 5.2.1

Software	Build Number	Details
VCF SDDC Manager Appliance (VCF-SDDC-Manager-Appliance-5.2.1.0-24307856.o va)	24307856	Only required for convert workflows.
VCF Import Tool (vcf-brownfield- import-5.2.1.2-24494579.tar. gz)	24494579	Required for convert, import, and sync workflows.
NOTE The 5.2.1.2 VCF Import Tool supports VMware Cloud Foundation 5.2.1 and 5.2.1.1.		
VMware Software Install Bundle - NSX_T_MANAGER 4.2.1.0 (bundle-133764.zip)	24304122	Required for convert or import workflows.

Table 184: Software Requirements for VMware Cloud Foundation 5.2.1.1

Software	Build Number	Details
VCF SDDC Manager Appliance (VCF-SDDC-Manager-Appliance-5.2.1.1-24397777.o va)	24397777	Only required for convert workflows.
VCF Import Tool (vcf-brownfield- import-5.2.1.2-24494579.tar. gz)	24494579	Required for convert, import, and sync workflows.

Table continued on next page

Continued from previous page

Software	Build Number	Details
NOTE The VCF Import Tool 5.2.1.1 has been replaced with the VCF Import Tool 5.2.1.2. Do not use the VCF Import Tool 5.2.1.1.		
VMware Software Install Bundle - NSX_T_MANAGER 4.2.1.0 (bundle-133764.zip)	24304122	Required for convert or import workflows.

VCF Import Tool Options and Parameters

You use the VCF Import Tool for converting or importing your existing vSphere environments into VMware Cloud Foundation.

The VCF Import Tool is distributed as a .tar package. Once extracted there is a python script called `vcf_brownfield.py`. The following options are available

Table 185: Script Actions

Action	Additional Information
-h, --help	Shows the VCF Import Tool help.
-v, --version	Displays the VCF Import Tool version.
convert	Converts existing vSphere infrastructure into the management domain in SDDC Manager.
check	Checks whether a vCenter is suitable to be imported into SDDC Manager as a workload domain.
import	Imports a vCenter as a VI workload domain into SDDC Manager.
sync	Syncs an imported VI workload domain or a VI workload domain deployed from SDDC Manager. See Manage Workload Domain Configuration Drift Between vCenter Server and SDDC Manager .
deploy-nsx	Deploys NSX Manager as a standalone operation. See Deploy NSX Manager for Workload Domains .
precheck	Runs prechecks on vCenter.

Table 186: Script Parameters

Parameters	Additional Information
--vcenter	Target vCenter Server for the current operation.
--sso-user	SSO administrator user for the target vCenter Server.
--sso-password	SSO administrator password for the target vCenter Server. Used for prevalidation only.

Table continued on next page

Continued from previous page

Parameters	Additional Information
--domain-name	Workload domain name to be assigned to the target environment during convert/import.
--nsx-deployment-spec-path	Absolute path to the NSX deployment spec json file.

Convert a vSphere Environment to a Management Domain or Import a vSphere Environment as a VI Workload Domain in VMware Cloud Foundation

You complete the necessary procedures to prepare your environment for converting or importing an existing vSphere environment to VMware Cloud Foundation and then validate the environment after the process is complete.

Convert to a Management Domain

Follow this procedure if your environment does not have SDDC Manager deployed and you want to convert an existing vSphere environment to the VMware Cloud Foundation management domain.

Review the following topics before proceeding:

- [Considerations Before Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#)
- [Network and Compute Requirements for Converting or Importing Existing vSphere Environments](#)
- [Download Software for Converting or Importing Existing vSphere Environments](#)

Step	Task Name	Additional Info
1	Copy the VCF Import Tool to the Target vCenter Appliance	
2	Run a Precheck on the Target vCenter Before Conversion	The precheck determines if the environment can be converted to the management domain
3	Remove the VCF Import Tool from vCenter.	The convert operation is run from SDDC Manager, once deployed.
4	Deploy the SDDC Manager Appliance on the Target vCenter	
5	Generate an NSX Deployment Specification for Converting or Importing Existing vSphere Environments	
6	Upload the Required Software to the SDDC Manager Appliance	
7	Run a Detailed Check on the Target vCenter Before Conversion or Import	
8	Convert or Import the vSphere Environment into the SDDC Manager Inventory	The workload domain will be marked type: MGMT
9	Add Licenses for Converted or Imported Workload Domains in SDDC Manager	

Table continued on next page

Continued from previous page

Step	Task Name	Additional Info
10	Validate a Converted Management Domain or Imported VI Workload Domain	

If you make changes to an imported or converted environment using the vSphere client, you can use the sync workflow in the VCF Import Tool to manually update the SDDC Manager inventory with any out-of-band changes applied from vCenter Server. See [Manage Workload Domain Configuration Drift Between vCenter Server and SDDC Manager](#).

Import as a VI Workload Domain

Follow this procedure if your environment already has SDDC Manager deployed and you want to import an existing vSphere environment as a VI workload domain.

You can convert/import up to a combined total of 24 workload domains.

Review the following topics before proceeding:

- [Considerations Before Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#)
- [Network and Compute Requirements for Converting or Importing Existing vSphere Environments](#)
- [Download Software for Converting or Importing Existing vSphere Environments](#)

Step	Task Name	Additional Info
1	Upload the Required Software to the SDDC Manager Appliance	
2	Generate an NSX Deployment Specification for Converting or Importing Existing vSphere Environments	
3	Run a Detailed Check on the Target vCenter Before Conversion or Import	
4	Convert or Import the vSphere Environment into the SDDC Manager Inventory.	The workload domain will be marked type: VI
5	Add Licenses for Converted or Imported Workload Domains in SDDC Manager	
6	Validate a Converted Management Domain or Imported VI Workload Domain	

If you make changes to an imported or converted environment using the vSphere client, you can use the sync workflow in the VCF Import Tool to manually update the SDDC Manager inventory with any out-of-band changes applied from vCenter Server. See [Manage Workload Domain Configuration Drift Between vCenter Server and SDDC Manager](#).

Copy the VCF Import Tool to the Target vCenter Appliance

You use the VCF Import Tool to run a precheck on the vCenter appliance that you are planning to convert to the VMware Cloud Foundation management domain.

1. Log in to the target vCenter Server as **root** at https://<vcenter_server_fqdn>:5480.

2. Navigate to **Access** and click **Edit** under **Access Settings**.
3. Switch on **Enable SSH Login** and click **OK**.
4. SSH to the vCenter Server as user **root**.
5. Enable Shell, if it is not already enabled.

```
shell
```

6. Change the default shell from `/bin/appliancesh` to `/bin/bash`.

```
chsh -s /bin/bash root
```

7. Create a directory for the VCF Import Tool. For example:

```
mkdir /tmp/vcfimport
```

8. Copy over the required software into the directory.

File name
vcf-brownfield-import-<buildnumber>.tar.gz

9. Extract the bundle.

```
tar -xvf vcf-brownfield-import-<buildnumber>.tar.gz
```

Run a Precheck on the Target vCenter Before Conversion

Before you perform a management domain convert operation, you must run a quick precheck to validate that the target vCenter configuration is supported for conversion.

This precheck is read-only. No modifications are made to the vCenter Server.

NOTE

This only applies to the management domain vCenter Server. Importing a VI workload domain runs a different set of checks.

1. SSH to the vCenter Server VM as user **root**.
2. Navigate to the directory where you copied the VCF Import Tool. For example:

```
cd /tmp/vcfimport/vcf-brownfield-import-<buildnumber>/vcf-brownfield-toolkit
```

3. Run the following command to precheck the target vCenter.

```
python3 vcf_brownfield.py precheck --vcenter '<my-vcenter-address>' --sso-user '<my-sso-username>' --sso-password '<my-sso-password>'
```

4. If any prechecks fail, refer to the troubleshooting section of this guide for more information.

See [VCF Import Tool Troubleshooting](#).

After the precheck succeeds, remove the VCF Import Tool from the vCenter appliance.

Deploy the SDDC Manager Appliance on the Target vCenter

Deploy the SDDC Manager appliance on the target vCenter before converting the vCenter to the VMware Cloud Foundation management domain.

You deploy the SDDC Manager appliance by using the OVA that you downloaded. See [Download Software for Converting or Importing Existing vSphere Environments](#).

1. In a web browser, log in to the target vCenter Server by using the vSphere Client.
`https://<vcenter_server_fqdn>/ui`
2. Select **Menu > VMs and Templates**.
3. In the inventory expand vCenter Server > Datacenter.
4. Right-click the management folder and select **Deploy OVF template**.
5. On the Select an OVF template page, select **Local file**, click **Upload files**, browse to the location of the SDDC Manager OVA file, click **Open**, and click **Next**.
6. On the Select a name and folder page, in the Virtual machine name text box, enter a virtual machine name, and click **Next**.
7. On the Select a compute resource page, click **Next**.
8. On the Review details page, review the settings and click **Next**.
9. On the License agreements page, accept the license agreement and click **Next**.
10. On the Select storage page, select the vSAN datastore and click **Next**.
11. On the Select networks page, from the Destination network drop-down menu, select the management network distributed port group and click **Next**.
12. On the Customize template page, enter the following values and click **Next**.

Setting	Description
Enter root user password	root user password
Enter login (vcf) user password	vcf user password
Enter basic auth user password	admin user password
Enter backup (backup) user password	backup user password
Enter Local user password	admin@local password
Hostname	FQDN for the appliance
NTP sources	The NTP server details for the appliance
Enable FIPS	Select to enable FIPS if the vCenter for the new management domain is FIPS-enabled
Network 1 IP address	The IP address for the appliance
Network 1 Subnet Mask	The subnet mask for the appliance
Network Default Gateway	The default gateway for the appliance
DNS Domain name	The domain name for the appliance
Domain search path	The domain search path(s) for the appliance
Domain name servers	The DNS servers for the appliance

13. On the Ready to complete page, click **Finish** and wait for the process to complete.
14. When the SDDC Manager appliance deployment completes, expand the management folder.
15. Right-click the SDDC Manager appliance and select **Power > Power On**.

Generate an NSX Deployment Specification for Converting or Importing Existing vSphere Environments

To deploy NSX Manager when you import or convert a vSphere environment into VMware Cloud Foundation, you must create an NSX deployment specification.

NSX deployment requires a minimum of 3 hosts.

When you deploy an NSX Manager cluster during a convert or import operation, it uses NSX-VLAN networking. See [Considerations Before Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#) for information about NSX-VLAN networking limitations.

1. Create a JSON file with the details of your NSX deployment.

For example:

```
{
  "license_key": "AAAAA-BBBBBB-CCCCC-DDDDD-EEEEEE",
  "form_factor": "medium",
  "admin_password": "*****",
  "install_bundle_path": "/nfs/vmware/vcf/nfs-mount/bundle/bundle-133764.zip",
  "cluster_ip": "172.16.11.71",
  "cluster_fqdn": "sfo-m01-nsx01.sfo.rainpole.io",
  "manager_specs": [{
    "fqdn": "sfo-m01-nsx01a.sfo.rainpole.io",
    "name": "sfo-m01-nsx01a",
    "ip_address": "172.16.11.72",
    "gateway": "172.16.11.1",
    "subnet_mask": "255.255.255.0"
  },
  {
    "fqdn": "sfo-m01-nsx01b.sfo.rainpole.io",
    "name": "sfo-m01-nsx01b",
    "ip_address": "172.16.11.73",
    "gateway": "172.16.11.1",
    "subnet_mask": "255.255.255.0"
  },
  {
    "fqdn": "sfo-m01-nsx01c.sfo.rainpole.io",
    "name": "sfo-m01-nsx01c",
    "ip_address": "172.16.11.74",
```

```
"gateway": "172.16.11.1",
"subnet_mask": "255.255.255.0"
}]
}
```

VCF version	Required NSX install bundle
5.2	bundle-124941.zip
5.2.1/5.2.1.1	bundle-133764.zip

2. Replace the content in the sample JSON with the information for your environment.
3. Copy the completed JSON file to the SDDC Manager appliance.

You will provide the path to the JSON file when you convert or import a vSphere environment to VMware Cloud Foundation.

Upload the Required Software to the SDDC Manager Appliance

Upload the VCF Import Tool bundle and the NSX deployment bundle to SDDC Manager appliance to enable converting or importing existing vSphere environments to VMware Cloud Foundation.

To use VCF Import Tool 5.2.1.2, delete any previous installations of the VCF Import Tool (typically in `/home/vcf/vcfimport`). Some of the steps for VCF Import Tool 5.2.1.2 are different than the steps for previous versions. These differences are called out in the procedure.

NOTE

The VCF Import Tool 5.2.1.1 has been replaced with the VCF Import Tool 5.2.1.2. Do not use the VCF Import Tool 5.2.1.1.

1. SSH to the SDDC Manager appliance as user **vcf**.
2. Copy the NSX deployment bundle `bundle-buildnumber.zip` to the `/nfs/vmware/vcf/nfs-mount/bundle/` folder.
3. Copy the VCF Import Tool to the SDDC Manager appliance.

- a) Create a folder for the VCF Import Tool.

```
mkdir /home/vcf/vcf-import-package
```

- b) Copy `vcf-brownfield-import-buildnumber.tar.gz` to the folder.

- c) Navigate to the folder and extract the bundle.

```
tar -xvf vcf-brownfield-import-buildnumber.tar.gz
```

- d) For VCF Import Tool 5.2.1.2, switch to the **root** account and run the install script.

```
su
```

```
cd vcf-brownfield-import-buildnumber
```

```
./install.sh
```

After the script completes successfully, switch to the **vcf** user and navigate to `/home/vcf/vcf-import-package/vcf-brownfield-import-version/vcf-brownfield-toolset`.

- e) Verify the scripts extracted correctly.

In the `/home/vcf/vcf-import-package/vcf-brownfield-import-<version>/vcf-brownfield-toolset` directory, run the following command:

```
python3 vcf_brownfield.py --help
```

This command should return the help information.

Run a Detailed Check on the Target vCenter Before Conversion or Import

Before you perform a management domain convert operation or a VI workload domain import operation, you must perform a detailed check to ensure that the existing vSphere environment's configuration is supported for convert or import.

1. SSH to the SDDC Manager appliance as user `vcf`.
2. Navigate to the directory where you copied the VCF Import Tool. For example:

```
cd /home/vcf/vcf-import-package/vcf-brownfield-import-<buildnumber>/vcf-brownfield-toolset
```

3. Run the following command to check that the vSphere environment can be converted or imported.

```
python3 vcf_brownfield.py check --vcenter '<my-vcenter-address>' --sso-user '<my-sso-username>'
```

4. If any checks fail, refer to the guardrails YAML file for information on the failed check. Refer to the troubleshooting section of this guide for more information on remediation.

See [VCF Import Tool Troubleshooting](#).

Convert or Import the vSphere Environment into the SDDC Manager Inventory

After you validate the target vCenter, you convert or import it into the SDDC Manager inventory.

1. SSH to the SDDC Manager VM as user `vcf`.
2. Navigate to the directory where you copied the VCF Import Tool. For example:

```
cd /home/vcf/vcf-import-package/vcf-brownfield-import-<buildnumber>/vcf-brownfield-toolset
```

3. Run the `vcf_brownfield.py` script and enter the required passwords when prompted.

Operation	Command	Additional Information
Convert an existing vSphere environment to a VCF management domain and deploy VLAN-backed NSX	<pre>python3 vcf_brownfield.py convert --vcenter '<vcenter-fqdn>' --sso-user '<sso-user>' --domain-name '<wld-domain-name>' --nsx- deployment-spec-path '<nsx- deployment-json-spec-path>'</pre>	Run once against the management domain vCenter Server
Import an existing vSphere environment as a VI workload	<pre>python3 vcf_brownfield.py import --vcenter '<vcenter- fqdn>' --sso-user '<sso-</pre>	Run once per VI Workload Domain

Table continued on next page

Continued from previous page

Operation	Command	Additional Information
domain and deploy VLAN-backed NSX	user>' --domain-name '<wld-domain-name>' --nsx-deployment-spec-path '<nsx-deployment-json-spec-path>'	

NOTE

During a convert operation, take a snapshot of the SDDC Manager appliance from vCenter when prompted.

- Inspect the command outputs highlighted in yellow. All should be status code 200.

NOTE

If any guardrails fail, refer to the guardrails YAML file for information on the failed guardrail. Refer to the troubleshooting section of this guide for more information on remediation. See [troubleshooting.dita](#). If the NSX deployment fails you can retry by following [deploy-nsx-manager-for-imported-workload-domains-or-workload-domains-with-vsphere-networking.dita](#).

- For VCF Import Tool 5.2 and 5.2.1 switch to the **root** account.

```
su -
```

- Restart all SDDC Manager services.

```
echo 'y' | /opt/vmware/vcf/operationsmanager/scripts/cli/sddcmanager_restart_services.sh
```

- Delete the snapshot of the SDDC Manager appliance.

Once all the SDDC Manager services have restarted the new workload domain (management domain or VI workload domain) should appear in the SDDC Manager UI.

Add Licenses for Converted or Imported Workload Domains in SDDC Manager

In order to be able to add clusters or VI workload domains to imported or converted environments, you must add your licenses to the SDDC Manager inventory.

- Log in to the SDDC Manager UI.
- In the navigation pane, click **Administration** > **Licensing**.
- Click **+ License Key**.
- Select a product from the drop-down menu.
- Enter the license key.
- Enter a description for the license.
- Click **Add**.
- Repeat for each license.

Validate a Converted Management Domain or Imported VI Workload Domain

After you successfully convert an existing environment to a management domain or import an environment as a VI workload domain, run an upgrade precheck to identify any potential issues.

1. Log in to the SDDC Manager UI.
2. Navigate to **Workload Domains** and click on the workload domain name.
3. Click the **Updates** tab and click **Run Precheck**.
4. Under **Target Version**, select **General Upgrade Readiness** and select all components.
5. Click **Run Precheck**.
6. Review the results.

VCF Import Tool Troubleshooting

This section describes troubleshooting issues and procedures for VCF Import Tool. It contains common troubleshooting scenarios encountered during convert or import processes and provides solutions for these scenarios.

General Troubleshooting for the VCF Import Tool

Identify and resolve common issues during the convert or import processes.

Review `/output/vcf_brownfield.log` for detailed information on any issues you encounter when using the VCF Import Tool to convert or import an existing vSphere environment.

Importing a VI Workload Domain Fails while Importing Trusted Root Certificates

While importing a VI Workload domain using the VCF Import Tool, you might encounter an error message when importing vCenter trusted root certificates.

Error message:

```
create trusted root chain failed : AFD Native Error Occured: 1006
```

This is caused by a high number of cert entries in the `TRUSTED_ROOT_CRLS` store. See <https://knowledge.broadcom.com/external/article?legacyId=70656> for information about how to resolve the issue. You can retry importing the VI workload domain after performing the steps in the KB.

VCF Import Tool Guardrail Troubleshooting

This section helps you identify, analyze, and resolve common configuration or integration issues when importing or converting existing vSphere environments in VMware Cloud Foundation using the VCF Import Tool.

ESX Upgrade Policy Guardrail Failure

When you use the VCF Import Tool to import or convert a vSphere environment, you may encounter a warning related to the ESX Upgrade Policy guardrail. This topic helps to resolve the warning.

In SDDC Manager, the default upgrade policy applies to all clusters, while in vSphere each cluster has a distinct upgrade policy. This can create a scenario where the ESX Upgrade Policy configured in vCenter does not match what SDDC Manager expects.

This issue causes a warning only. You can proceed with the convert or import process without remediating the issue.

1. Update the ESX upgrade policy in vCenter to match what SDDC Manager expects as described in the following table. See [How to Configure the vSphere Lifecycle Manager Remediation Settings](#) for more information about how to change the cluster settings in the vSphere Client.

Table 187: Default ESX upgrade policy in SDDC Manager

Upgrade Policy	vLCM Images	vLCM Baselines
Quick Boot	Enabled	
VM power state	Do not change power state	
VM migration	Migrate powered off and suspended VMs to other hosts in the cluster	
Retry Policy		
Failure action	RETRY	FailTask
Retry delay	300	
Retry count	3	0
HA admission control	Enabled	Disabled
Distributed Power Management	Disabled	
Hardware compatibility issues	Disabled	
Quick Patch	Disabled	

Troubleshooting Guardrail Issues Requiring Manual SDDC Manager Database Updates

When you are using the VCF Import Tool, certain guardrail messages may require you to perform a manual update to the SDDC Manager database. If you encounter any of these guardrail issues, modify the example commands to resolve the issues.

1. SSH to SDDC Manager as **vcf**.
2. Switch to the **root** account.

```
su
```

3. Run the command to remediate the issue for the specific guardrail.

Replace the values in the example commands with the information for you environment.

Guardrail	Action	Example Command
vCenter version patch level could not be detected, defaulting to 00000 in SDDC Manager inventory.	Update the vCenter <i>version</i> field in the SDDC Manager database with the correct patch level.	<pre>psql -h localhost -U postgres -d platform -c "update vcenter set version='8.0.3.00000-23514763' where vm_hostname='sfo-m01-vc01.sfo.rainpole.io'"</pre>
Could not find version and build number for ESXi: {}.	Please update the host <i>version</i> field in the SDDC Manager database.	<pre>psql -h localhost -U postgres -d platform -c "update host set version='8.0.3-23637092' where hostname='sfo01-m01-esx01.sfo.rainpole.io'"</pre>
Could not find vCenter VM name.	Update the vCenter <i>vmName</i> field in the SDDC Manager database .	<pre>psql -h localhost -U postgres -d platform -c "update vcenter set vm_name='sfo-m01-vc01'"</pre>

Table continued on next page

Continued from previous page

Guardrail	Action	Example Command
		<code>where vm_hostname='sfo-m01-vcf01.sfo.rainpole.io'"</code>
Could not find SDDC Manager VM name.	Update the SDDC Manager Controller <i>vmName</i> field in the SDDC Manager database.	<code>psql -h localhost -U postgres -d platform -c "update sddc_manager_controller set vm_name='sfo-vcf01' where vm_hostname='sfo-vcf01.sfo.rainpole.io'"</code>
Could not find IP address for SDDC Manager: {}.	Update the SDDC Manager Controller <i>vmManagementIpAddress</i> field in the SDDC Manager database .	<code>psql -h localhost -U postgres -d platform -c "update sddc_manager_controller set vm_management_ip_address='172.16.11.59' where vm_hostname='sfo-vcf01.sfo.rainpole.io'"</code>
Could not read datastore from VM datastore path: {}.	Update the vCenter and PSC <i>datastoreName</i> field in the SDDC Manager database.	<code>psql -h localhost -U postgres -d platform -c "update vcenter set datastore_name='sfo-m01-ds01' where vm_hostname='sfo-m01-vcf01.sfo.rainpole.io'"</code>
Could not find datastore for VM: {}.	Update the vCenter and PSC <i>datastoreName</i> field in the SDDC Manager database .	<code>psql -h localhost -U postgres -d platform -c "update vcenter set datastore_name='sfo-m01-ds01' where vm_hostname='sfo-m01-vcf01.sfo.rainpole.io'"</code>
Could not find gateway IP for ESXi: {}.	Update the host <i>gateway</i> field in the SDDC Manager database .	<code>psql -h localhost -U postgres -d platform -c "update host set gateway='172.16.11.1' where hostname='sfo01-m01-esx01.sfo.rainpole.io'"</code>
Could not find the management vmKernel for ESXi: {}. Skipping population of host <i>management_ip_address</i> and <i>subnet</i> fields.	Update the <i>management_ip_address</i> and <i>subnet</i> fields in the SDDC Manager database.	<code>psql -h localhost -U postgres -d platform -c "update host set management_ip_address='172.16.11.101' where hostname='sfo01-m01-esx01.sfo.rainpole.io'"</code>

Table continued on next page

Continued from previous page

Guardrail	Action	Example Command
		<pre>psql -h localhost -U postgres -d platform -c "update host set subnet='255.255.255.0' where hostname='sfo01-m01- esx01.sfo.rainpole.io'"</pre>
Could not find the vMotion vmKernel for ESXi: {}. Skipping population of host <i>vmotion_ip_address</i> field.	Update the <i>vmotion_ip_address</i> field in the SDDC Manager database.	<pre>psql -h localhost -U postgres -d platform -c "update host set vmotion_ip_address=' 172.16.12.101' where hostname='sfo01-m01- esx01.sfo.rainpole.io'"</pre>
Could not find managed object ID of ESXi: {}. Skipping population of ESXi <i>source_id</i> field.	Update the <i>source_id</i> field in the SDDC Manager database.	<pre>psql -h localhost -U postgres -d platform -c "update host set source_id='host-25' where hostname='sfo01-m01- esx01.sfo.rainpole.io'"</pre>
Could not find vSAN FTT configuration of cluster: {}.	Update the <i>ftt</i> field for this cluster in the SDDC Manager database.	<pre>psql -h localhost -U postgres -d platform -c "update cluster set ftt='1' where name='sfo-m01-cl01'"</pre>
Could not identify primary datastore for cluster: {}.	Update the <i>primaryDatastoreName</i> and <i>primaryDatastoreType</i> fields for the cluster in the SDDC Manager database.	<pre>psql -h localhost -U postgres -d platform -c "update cluster set primary_datastore_name=' sfo-m01-cl01-ds-vsan01' where name='sfo-m01-cl01'"</pre>
		<pre>psql -h localhost -U postgres -d platform -c "update cluster set primary_datastore_type='VSA N' where name='sfo-m01- cl01'"</pre>
Could not find the subdomain name from SDDC Manager hostname: {}.	Update the <i>subDomain</i> field in <i>dnsInfo</i> of <i>SystemInfo</i> in the SDDC Manager database .	<pre>psql -h localhost -U postgres -d platform -c "update SystemInfo set dns_info= jsonb_set(dns_info, '{subDomain}', '"sfo.rainpole.io"', true)"</pre>

Table continued on next page

Continued from previous page

Guardrail	Action	Example Command
Could not find root domain from SDDC Manager subdomain name: {}.	Update the <i>rootDomain</i> field in <i>dnsInfo</i> of <i>SystemInfo</i> in the SDDC Manager database .	<pre>psql -h localhost -U postgres -d platform -c "update SystemInfo set dns_info=jsonb_set(dns_info, '{rootDomain}', '"rainpole.io"', true)"</pre>
Could not find DNS configuration of SDDC Manager.	Update the <i>primaryDns</i> and <i>secondaryDns</i> fields in <i>dnsInfo</i> of <i>SystemInfo</i> in the SDDC Manager database .	<pre>psql -h localhost -U postgres -d platform -c "update SystemInfo set dns_info=jsonb_set(dns_info, '{primaryDns}', '"172.16.11.4"', true)"</pre>
Could not find NTP configuration of SDDC Manager.	Update <i>ntpS</i> field in <i>ntpInfo</i> of <i>SystemInfo</i> in the SDDC Manager database .	<pre>psql -h localhost -U postgres -d platform -c "update SystemInfo set ntp_info=jsonb_set(ntp_info, '{ntpS}', '"ntp0.sfo.rainpole.io", "ntp1.sfo.rainpole.io"', true)"</pre>
Could not find SDDC Manager version.	Verify that the <code>/opt/vmware/vcf/version.txt</code> file has read permissions for user <code>vcf</code> . Update the SDDC Manager Controller <i>version</i> field in the SDDC Manager database.	<pre>psql -h localhost -U postgres -d platform -c "update sddc_manager_controller set version='5.2.0.0-23684695'"</pre>

Managing Workload Domains in VMware Cloud Foundation

Workload domains are logical units that carve up the compute, network, and storage resources of the VMware Cloud Foundation system. The logical units are groups of ESXi hosts managed by vCenter Server instances with specific characteristics for redundancy and VMware best practices.

Each workload domain include these VMware capabilities by default:

- vCenter Server Appliance
- vSphere High Availability (HA)
- vSphere Distributed Resource Scheduler (DRS)
- vSphere Distributed Switch
- Shared Storage (vSAN, vSAN ESA, NFS, vVols, Fibre Channel)

NOTE

The management domain default cluster datastore must use vSAN.

- NSX Manager Cluster

See the [VMware Private AI Foundation with NVIDIA Guide](#) for an overview of the components of VMware Private AI Foundation with NVIDIA and high-level workflows for development and production use cases.

NOTE

VMware Private AI Foundation with NVIDIA is supported with VMware Cloud Foundation 5.1.1 and later.

About VI Workload Domains

When deploying a VI workload domain, you specify the storage, name, compute, and networking details. Based on the selected storage, you provide vSAN parameters, NFS share details, VMFS on FC datastore information, or vVols storage details. You then select the hosts for the VI workload domain and start the workflow.

When you deploy a new VI workload domain, VMware Cloud Foundation deploys a new vCenter Server for that workload domain. The vCenter Server is associated with a vCenter Single Sign-On Domain (SSO) to determine the local authentication space. Prior to VMware Cloud Foundation 5.0, the management vCenter Server and all VI workload domain vCenter Servers were members of a single vSphere SSO domain, joined together with vCenter Enhanced Linked Mode. Starting with VMware Cloud Foundation 5.0, when you deploy a new VI workload domain, you can choose to join the management domain SSO domain, or create a new SSO domain.

The workflow automatically:

- Deploys a vCenter Server Appliance for the new VI workload domain within the management domain. By using a separate vCenter Server instance per VI workload domain, software updates can be applied without impacting other VI workload domains. It also allows for each VI workload domain to have additional isolation as needed.
- Connects the specified ESXi servers to this vCenter Server instance and groups them into a cluster. Each host is configured with the port groups applicable for the VI workload domain.
- Configures networking on each host.
- Configures vSAN, NFS, VMFS on FC, or vVols storage on the ESXi hosts.
- For the first VI workload domain, the workflow deploys a cluster of three NSX Managers in the management domain and configures a virtual IP (VIP) address for the NSX Manager cluster. The workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for high availability. Subsequent VI workload domains can share an existing NSX Manager cluster or deploy a new one. To share an NSX Manager cluster, the VI workload domains must use the same update method. The VI workload domains must both use vSphere Lifecycle Manager (vLCM) images, or they must both use vLCM baselines.
- By default, VI workload domains do not include any NSX Edge clusters and are isolated. To provide north-south routing and network services, add one or more NSX Edge clusters to a VI workload domain. See [Managing NSX Edge Clusters in VMware Cloud Foundation](#).

NOTE

Starting with VMware Cloud Foundation 5.2, when you deploy a new VI workload domain, it uses the same versions of vCenter Server and NSX Manager that the management domain uses. For example, if you applied an async patch to the vCenter Server in the management domain, a new VI workload domain will deploy the same patched version of vCenter Server.

VMware Cloud Foundation 5.2 and later support importing existing vSphere environments as VI workload domains. See [Converting or Importing Existing vSphere Environments into VMware Cloud Foundation](#) for more information.

Prerequisites for a Workload Domain

Review the prerequisites before you deploy a VI workload domain.

- If you plan to use DHCP for the NSX host overlay network, a DHCP server must be configured on the NSX host overlay VLAN for the VI workload domain. When VMware NSX creates NSX Edge tunnel endpoints (TEPs) for the VI workload domain, they are assigned IP addresses from the DHCP server.

NOTE

If you do not plan to use DHCP, you can use a static IP pool for the NSX host overlay network. The static IP pool is created or selected as part of VI workload domain creation.

- A minimum of three hosts marked with the appropriate storage must be available in your SDDC Manager inventory.

NOTE

If you are using NFS, VMFS on FC, or vVols as principal storage, and the VI workload domain is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- To create a VI workload domain with NFS storage, the hosts must be commissioned with NFS as the storage type and must be associated with an NFS network pool.
- To create a VI workload domain with vSAN storage, the hosts must be commissioned with vSAN as the storage type and must be associated with a vSAN network pool.
- To create a VI workload domain with VMFS on FC storage, the hosts must be commissioned with VMFS on FC as the storage type and must be associated with a vMotion only or vMotion and NFS network pool.
- To create a VI workload domain with vVols on FC, the hosts must be commissioned with vVols on FC as the storage protocol type and must be associated with vMotion only or vMotion and NFS.
- To create a VI workload domain with vVols on iSCSI, the hosts must be commissioned with vVols on iSCSI as the storage protocol type and must be associated with vMotion and iSCSI.
- To create a VI workload domain with vVols on NFS, the hosts must be commissioned with vVols on NFS as the storage protocol type and must be associated with vMotion and NFS.

NOTE

vVols requires that a VASA provider must be added to the SDDC Manager inventory to create a VI workload domain with vVols storage. See [vVols Storage with VMware Cloud Foundation](#).

For information on adding hosts to your inventory, see [Managing ESXi Hosts in VMware Cloud Foundation](#).

- There must be a free pNIC on each host to be used for the VI workload domain.
- To create a VI workload domain that uses vSphere Lifecycle Manager to apply cluster images to all hosts in the cluster, you must have a cluster image available. See [Managing vSphere Lifecycle Manager Images in VMware Cloud Foundation](#).
- The install bundles for the versions of NSX Manager and vCenter Server that are running in the management domain must be available in SDDC Manager before you can create a VI workload domain. For example, if you have patched the versions of NSX Manager and/or vCenter Server in the management domain to a version higher than what is listed in the BOM, you must download the new install bundles. See [Downloading Install Bundles for VMware Cloud Foundation](#). You can refer to <https://knowledge.broadcom.com/external/article?legacyId=88287> for information about the install bundles required for specific async patches.
- Decide on a name for your VI workload domain. Each VI workload domain must have a unique name. It is good practice to include the region and site information in the name because resource object names (such as host and vCenter names) are generated based on the VI workload domain name. The name can be three to 20 characters long and can contain any combination of the following:
 - Lowercase alphabetic characters
 - Uppercase alphabetic characters
 - Numbers

NOTE

Spaces are not allowed in any of the names you specify when creating a VI workload domain.

- Decide on the following passwords:
 - vCenter Server root password
 - NSX Manager admin password

Although the individual VMware Cloud Foundation components support different password requirements, you must set passwords following a common set of requirements across all components:

- Minimum length: 12
- Maximum length: 16
- At least one lowercase letter, one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ *
- Must NOT include:

- A dictionary word
 - A palindrome
 - More than four monotonic character sequences
 - Three of the same consecutive characters
- Verify that you have the completed Planning and Preparation Workbook with the VI workload domain deployment option included.
 - The IP addresses and Fully Qualified Domain Names (FQDNs) for the vCenter Server and NSX Manager instances must be resolvable by DNS.
 - If you are using NFS storage for the VI workload domain, you need the following information:
 - Datastore name
 - Path to the NFS share
 - IP address of the NFS server
 The NFS share and server must be accessible from the management network. You must have read/write permission to the NFS share.
 - If you are using VMFS on FC storage for the VI workload domain, you must configure zoning, mount the associated volumes and create the datastore on the hosts.
 - To use the **License Now** option, you must have valid license keys for the following products:
 - VMware NSX
 - vSAN (No license required for NFS or VMFS on FC)
 - vSphere
 Because vSphere and vSAN licenses are per CPU, ensure that you have sufficient licenses for the ESXi hosts to be used for the VI workload domain. See [Managing License Keys in](#) .
 - If you plan to deploy a VI workload domain that has its vSphere cluster at a remote location, you must meet the following requirements:
 - Dedicated WAN connectivity is required between central site and remote site.
 - Primary and secondary active WAN links are recommended for connectivity from the central site to the remote site. The absence of WAN links can lead to two-failure states, WAN link failure, or NSX Edge node failure, which can result in unrecoverable VMs and application failure at the remote site.
 - Minimum bandwidth of 10 Mbps and latency of 100 ms is required between the central site and remote site. The network at the remote site must be able to reach the management network at the central site. DNS and NTP server must be available locally at or reachable from the remote site.
 - See [VMware Configuration Maximums](#) for limitations related to VI workload domains at remote locations.
 - See [VMware Cloud Foundation Edge Design Considerations](#) for more information about design options for deploying scalable edge solutions.
 - If you plan to deploy a multi-rack compute VI workload domain, you must meet the following requirements:
 - A dedicated host uplink profile for each rack with a separate host TEP transport VLAN ID per rack.
 - A dedicated host TEP IP pool for each rack, with a subnet allocated per rack and a gateway for the subnet.
 - A dedicated NSX host sub-transport node profile for each rack.
 - A dedicated network pool in SDDC Manager for each rack.
 - A dedicated vSAN fault domain for each rack.

Deploy NSX Manager for Workload Domains

If you have a workload domain without NSX Manager, you can add an NSX Manager cluster with NSX-VLAN networking to the workload domain using the VCF Import Tool.

Ensure that the install bundle for a supported version of NSX is available in SDDC Manager (**Lifecycle Management** > **Bundle Management**) or the installer zip file is available on the SDDC Manager appliance. You can download the installer zip file (`bundle-124941.zip` for VCF 5.2; `bundle-133764.zip` for VCF 5.2.1) from the [Broadcom Support portal](#).

Identify the FQDN of the vCenter Server for the workload domain where you want to deploy NSX Manager.

The workload domain to which you are deploying NSX Manager must include a cluster with 3 or more hosts.

When you add an NSX Manager cluster using this procedure, it uses NSX-VLAN networking, not NSX-overlay networking. VI workload domains with NSX-VLAN networking do not support:

- Stretch clusters
 - Edge cluster
 - Application Virtual Networks (AVNs)
 - VMware Aria Suite Lifecycle
 - VMware Cloud Foundation with VMware Tanzu (Workload Management)
- Avi Load Balancer is supported in VI workload domains with NSX-VLAN networking.

1. Download the VCF Import Tool and copy it to the SDDC Manager appliance.

NOTE

Make sure to download the version of the VCF Import Tool for your version of VMware Cloud Foundation. There are different versions of the VCF Import Tool for VCF 5.2 and VCF 5.2.1.

You can skip this step if the VCF Import Tool already exists on the SDDC Manager appliance.

- a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
- b) Click **VMware Cloud Foundation 5.2**.
- c) Click **Drivers & Tools**.
- d) Click the download icon for the VCF Import Tool.
- e) SSH in to the SDDC Manager appliance using the `vcf` user account.
- f) Create a folder for import scripts.

```
mkdir /home/vcf/vcfimport
```

- g) Copy `vcf-brownfield-import-<version-number>.tar.gz` to `/home/vcf/vcfimport`.
- h) Navigate to the scripts directory and extract the scripts bundle.

```
cd /home/vcf/vcfimport
```

```
tar -xvf vcf-brownfield-import-<version-number>.tar.gz
```

- i) Browse to the `vcf-brownfield-toolset` directory and run:

```
python3 vcf_brownfield.py --help
```

2. Create an NSX deployment spec JSON and copy it to the SDDC Manager appliance.

For example:

```
{
  "license_key": "AAAAA-BBBBBB-CCCCC-DDDDD-EEEEEE",
  "form_factor": "medium",
  "admin_password": "<password>",
  "install_bundle_path": "/nfs/vmware/vcf/nfs-mount/bundle/bundle-124941.zip",
  "cluster_ip": "172.16.11.71",
  "cluster_fqdn": "sfo-m01-nsx01.sfo.rainpole.io",
```

```

"manager_specs": [{
  "fqdn": "sfo-m01-nsx01a.sfo.rainpole.io",
  "name": "sfo-m01-nsx01a",
  "ip_address": "172.16.11.72",
  "gateway": "172.16.11.1",
  "subnet_mask": "255.255.255.0"
},
{
  "fqdn": "sfo-m01-nsx01b.sfo.rainpole.io",
  "name": "sfo-m01-nsx01b",
  "ip_address": "172.16.11.73",
  "gateway": "172.16.11.1",
  "subnet_mask": "255.255.255.0"
},
{
  "fqdn": "sfo-m01-nsx01c.sfo.rainpole.io",
  "name": "sfo-m01-nsx01c",
  "ip_address": "172.16.11.74",
  "gateway": "172.16.11.1",
  "subnet_mask": "255.255.255.0"
}
}]
}

```

If the NSX install bundle is already available in SDDC Manager, you can omit the "install_bundle_path".

VCF version	Required NSX install bundle
5.2	bundle-124941.zip
5.2.1	bundle-133764.zip

3. Run the script to deploy NSX Manager.

- a) SSH in to the SDDC Manager appliance using the `vcf` user account.
- b) Browse to the `vcf-brownfield-toolset` directory and run:

```
python3 vcf_brownfield.py deploy-nsx --vcenter <vcenter-server-fqdn> --nsx-
deployment-spec-path <nsx-deployment-json-spec-path>
```

Delete a VI Workload Domain

You can delete a VI workload domain from SDDC Manager UI.

- If remote vSAN datastores are mounted on a cluster in the VI workload domain, then the VI workload domain cannot be deleted. To delete such VI workload domains, you must first migrate any virtual machines from the remote datastore to the local datastore and then unmount the remote vSAN datastores from vCenter Server.
- If you require access after deleting a VI workload domain, back up the data. The datastores on the VI workload domain are destroyed when it is deleted.
- Migrate the virtual machines that you want to keep to another workload domain using cross vCenter vMotion.
- Delete any workload virtual machines created outside VMware Cloud Foundation before deleting the VI workload domain.
- Delete any NSX Edge clusters hosted on the VI workload domain. See [KB 78635](#).

When you delete a workload domain, the clusters within that workload domain are deleted and the hosts are returned to the free pool with a `need_cleanup` host state.

Deleting a VI workload domain also removes the components associated with the VI workload domain from the management domain. This includes the vCenter Server instance and the NSX Manager cluster instances.

NOTE

If the NSX Manager cluster is shared with any other VI workload domains, it will not be deleted.

The network pools used by the workload domain are not deleted as part of the VI workload domain deletion process and must be deleted separately.

CAUTION

Deleting a workload domain is an irreversible operation. All clusters and virtual machines within the VI workload domain are deleted and the underlying datastores are destroyed.

It can take up to 20 minutes for a VI workload domain to be deleted. During this process, you cannot perform any operations on workload domains.

1. In the navigation pane, click **Inventory > Workload Domains**.
2. Click the vertical ellipsis (three dots) next to the VI workload domain you want to delete and click **Delete Domain**.

Add Cluster

Add Edge Cluster

Delete Domain

Rename Domain

3. On the **Delete Workload Domain** dialog box, click **Delete Workload Domain**.

A message indicating that the VI workload domain is being deleted appears. When the removal process is complete, the VI workload domain is removed from the domains table.

If you delete an isolated VI workload domain that created an NSX Manager cluster that is shared with another isolated VI workload domain, you need to register NSX Manager as a relying partner to the remaining VI workload domain. See <https://kb.vmware.com/s/article/95445>.


Decommission the hosts that were part of the VI workload domain, then re-image with ESXi and commission them again. See [Managing ESXi Hosts in VMware Cloud Foundation](#).

View Workload Domain Details

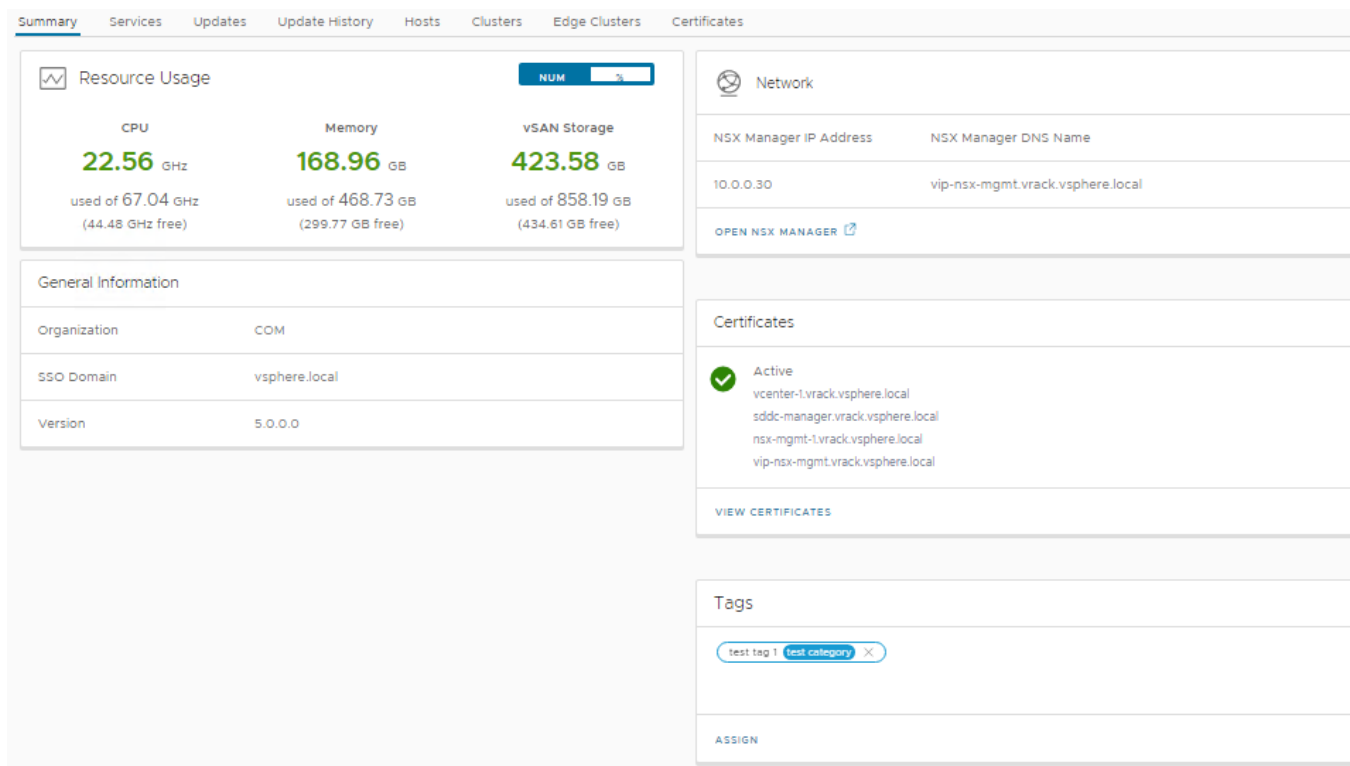
The Workload Domains page displays high level information about the workload domains in a VMware Cloud Foundation instance. CPU, memory, and storage utilized by the workload domain is also displayed here.

1. In the navigation pane, click **Inventory > Workload Domains**.

TIP

Click the show or hide columns icon  to view additional information about the workload domains, including the SSO domain.

2. In the workload domains table, click the name of the workload domain.



Tab	Information Displayed
Summary	Provides information about: <ul style="list-style-type: none"> • Resource usage: CPU, memory, and storage resources for the workload domain. • Network: NSX Manager IP address and DNS name. • General information, including SSO domain. • Certificates • Tags
Services	SDDC software stack components deployed for the workload domain's virtual environment and their IP addresses. Click a component name to navigate to that aspect of the virtual environment. For example, click vCenter Server to reach the vSphere Client for that workload domain. All the capabilities of a VMware SDDC are available to you in the VI workload domain's environment, such as creating, provisioning, and deploying virtual machines, configuring the software-defined networking features, and so on.
Updates	Available updates for the workload domain.
Update History	Updates applied to this workload domain.

Table continued on next page

Continued from previous page

Tab	Information Displayed
Hosts	Names, IP addresses, status, associated clusters, and capacity utilization of the hosts in the workload domain and the network pool they are associated with.
Clusters	Names of the clusters, number of hosts in the clusters, and their capacity utilization.
Edge Clusters	Names of the NSX Edge clusters, NSX Edge nodes, and their status.
Certificates	Default certificates for the VMware Cloud Foundation components. For more information, see Managing Certificates in VMware Cloud Foundation .

Expand a Workload Domain

You can expand the management domain or a VI workload domain to add resources to support additional workloads or availability.

To expand a domain, you can:

- Add a host from the SDDC Manager inventory to a vSphere cluster.
By adding an individual host to an existing workload domain, you can expand the amount of resources contained within an existing vSphere cluster up to the supported vSphere maximums.
- Add a new vSphere cluster to a workload domain.
Workload domains support multiple vSphere clusters. You can add an additional vSphere cluster to an existing workload domain to provide for increased capacity.

Add a Host to a vSphere Cluster Using the SDDC Manager UI

You can add an ESXi host to a vSphere cluster using the SDDC Manager UI. Adding an ESXi host to a vSphere cluster in a workload domain increases the available resources. You can add multiple hosts at the same time.


- Verify that a host is available in the SDDC Manager inventory. For information on commissioning hosts, see [Commission Hosts](#).
- Verify that the host you want to add is in an active state. See [View Host Inventory](#).
- If you choose **License Now**, you must have a valid vSphere license available in the SDDC Manager inventory. See [Add a Component License Key in the](#)
- Verify that the host to be added matches the configuration of the hosts already in the vSphere cluster. This allows the vSphere cluster configuration to remain balanced. If the host to be added does not match the pre-existing hosts in the vSphere cluster, the cluster will be unbalanced and a warning is displayed. The warning does not prevent the expansion and can be dismissed if needed.
- Verify that the host you are adding has the same type of principal storage as the existing hosts in the vSphere cluster. For the management domain, the host must use vSAN for principal storage. For VI workload domains, the host can use vSAN, NFS, VMFS on FC, or vVols for principal storage. A host using NFS for principal storage will automatically use the same NFS configuration as the other hosts in the vSphere cluster. For a host using VMFS on FC, you must configure zoning, mount the associated volumes, and create the datastore on the host before adding the host to a vSphere cluster. A host using vVols for principal storage will automatically use the same vVols configuration as the other hosts in the vSphere cluster.
- If the vSphere cluster hosts an NSX Edge cluster, you can only add new hosts with the same management, uplink, NSX Host Overlay, and NSX Edge Overlay networks (L2 uniform) as the existing hosts.
- If the vSphere cluster to which you are adding hosts uses a static IP pool for the NSX Host Overlay Network TEPs, that pool must include enough IP addresses for the hosts you are adding.
- Different NIC enumeration is allowed.

When you use the SDDC Manager UI to add an ESXi host to a vSphere cluster, it uses the cluster's existing vDS for management, vMotion, vSAN, and host overlay traffic.

NOTE

You can add hosts to or remove hosts from multiple different vSphere clusters in parallel. For example, you add three hosts to Cluster A, and while that task is running, you can start a separate task to add (or remove) four hosts to Cluster B. See [VMware Configuration Maximums](#) for information about the maximum number of add/remove hosts tasks that you can run in parallel.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the workload domains table, click the name of the workload domain that you want to expand.
3. Click the **Clusters** tab.
4. Click the name of the vSphere cluster where you want to add a host.

[Add Host](#)
[Mount Remote Datastore](#)
[Rename Cluster](#)
[Open in vSphere Client](#) 

5. Click **Actions** › **Add Host**.

6. Select the cluster expansion type.

This option only appears if the vSphere cluster hosts an NSX Edge cluster.

L2 Uniform	Select if all hosts you are adding to the vSphere cluster have the same management, uplink, NSX Overlay Host TEP, and NSX Edge Overlay Edge TEP networks as the existing hosts in the vSphere cluster.
L2 non-uniform and L3	Select if any of the hosts you are adding to the vSphere cluster have different networks than the existing hosts in the vSphere cluster. <p>IMPORTANT VMware Cloud Foundation does not support adding hosts to L2 non-uniform and L3 vSphere clusters that host an NSX Edge cluster.</p>

7. Select the host you want to add to the vSphere cluster and click **Next**.

NOTE

If the vSphere cluster includes DPU-backed hosts, you can only add DPU-backed hosts from the same vendor.

After selecting the host, you can choose vmnics to uplinks of the vDS. The enumeration of physical nics does not need to be the same. You can use any free vnic on the host.

8. On the Switch Configuration page, you can view the existing networking configuration on the cluster. It shows both vDS and NSX networking. Each vDS has two uplinks in it, and you can assign any free vmnics on the host to the existing uplinks on the vDS.
9. On the Licenses page, choose a licensing option.

Option	Description
License Now	Select the vSphere license to apply to the hosts.
License Later	ESXi hosts are deployed in evaluation mode. IMPORTANT After the hosts are added, you must switch to licensed mode by: <ul style="list-style-type: none"> • Adding a component license key in the SDDC Manager UI. See Add a Component License Key in the SDDC Manager UI. Or, • Adding a solution license key in the vSphere Client. See Managing vSphere Licenses.

10. Click **Next**.
11. Review the host, switch configuration, and license details and click **Finish**.

A message indicating that the host is being added will appear.

Add a vSphere Cluster to a Workload Domain Using the SDDC Manager UI

If you want to add capacity to a workload domain, you can use the SDDC Manager UI to add a vSphere cluster.

- Verify that there are at least three hosts available in the SDDC Manager inventory. For information on commissioning hosts, see [Commission Hosts](#).

NOTE

If the vSphere cluster is using NFS, VMFS on FC, or vVols as principal storage and is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- Ensure that the hosts you want to add to the vSphere cluster are in an active state.
- If you choose **License Now**, you must have valid license keys for vSphere and vSAN (if applicable) available in the SDDC Manager inventory. For more information, see [Add a Component License Key in the SDDC Manager UI](#).
- If you are using DHCP for the NSX Host Overlay Network, a DHCP server must be configured on the NSX Host Overlay VLAN of the management domain. When NSX creates TEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.

VMware Cloud Foundation 5.2.1 supports both vSphere Lifecycle Manager baseline and vSphere Lifecycle Manager image based clusters in the same workload domain. You can add a cluster that uses either vSphere Lifecycle Manager update method. For VMware Cloud Foundation 5.2, the vSphere Lifecycle Manager update method is specified at the workload domain level. When you add a new cluster to a workload domain, it uses the same update method as the other clusters in that workload domain.

If a workload domain has multiple clusters, each cluster can use a different type of principal storage, as long as all hosts within a vSphere cluster use the same type.

You can run multiple add cluster tasks at the same time.

To add a cluster that contains ESXi hosts with two data processing units (DPU), you must configure a custom vSphere Distributed Switch (VDS) with the following settings: :

- A single VDS that uses all four DPU-backed nics
- Uplinks (uplink1 through uplink4) are mapped to the DPU-backed nics
- The NSX network operational mode is set to **Enhanced Datapath - Standard**
- In the NSX uplink profile, uplink-1 and uplink-2 are Active and uplink-3 and uplink-4 are Standby

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. Click the vertical ellipsis (three dots) next to the workload domain to which you want to add a vSphere cluster and click **Add Cluster**.
3. Select the storage type for the vSphere cluster and click **Begin**.

For VMware Cloud Foundation 5.2, you can only activate vSAN ESA if the workload domain to which you are adding the cluster uses vSphere Lifecycle Manager images.

NOTE

vSAN Max requires vSAN ESA.

4. Enter a name for the vSphere cluster.
5. For VMware Cloud Foundation 5.2.1, select **Manage this cluster using vLCM images** to use vSphere Lifecycle Manager images as the update method.

If you do not select **Manage this cluster using vLCM images**, then the cluster uses vLCM baselines.

NOTE

- You must use vSphere Lifecycle Manager images for vSAN ESA and vSAN Max clusters.
- You must use vSphere Lifecycle Manager images in order to create vSphere clusters with only two hosts and those clusters must use NFS, VMFS on FC, or vVols as principal storage.
- Network offloading using SmartNICs (DPU-backed hosts) is not supported with vSphere Lifecycle Manager baselines.

6. Select a cluster image from the drop-down menu.

The option is only available for clusters that use vSphere Lifecycle Manager images.

NOTE

If the cluster image contains a different version of a vendor add-on or component than what is installed on the ESXi hosts you add to the cluster, the hosts will be remediated to use the cluster image during cluster creation.

7. Click **Next**.
8. If you selected vSAN storage for the vSphere cluster, the vSAN parameters page appears. The vSAN storage options are different for vSAN OSA and vSAN ESA.
 - a) For vSAN OSA, select the vSAN Storage Type:

vSAN HCI	<p>Provides storage and compute resources. Specify the level of availability you want configured for this cluster. The specified Failures To Tolerate (FTT) value determines the number of hosts required in the cluster.</p> <p>Select the check box to enable vSAN deduplication and compression.</p> <p>Click Next.</p>
vSAN Compute Cluster	<p>Provides compute resources only. Click Next and choose the remote datastore to provide storage to the new cluster.</p>

- b) For vSAN ESA, select the vSAN Storage Type:

vSAN HCI	<p>Provides storage and compute resources. SDDC Manager selects the following settings, which cannot be edited:</p> <ul style="list-style-type: none"> • Storage Type: Local vSAN datastore. • Storage Policy: Auto-policy management. <p>NOTE Based on the type of cluster and number of hosts, vSAN creates and assigns a default datastore policy for best capacity utilization after the cluster configuration is completed. Policy details can be viewed in the vSphere Client (Policies and Profiles › VM Storage Policies).</p> <p>Click Next.</p>
vSAN Max	<p>Provides storage resources only. You can mount a vSAN Max datastore on other vSAN ESA or vSAN computer clusters.</p> <p>SDDC Manager selects the following settings, which cannot be edited:</p> <ul style="list-style-type: none"> • Storage Type: Local vSAN datastore. • Storage Policy: Auto-policy management. <p>NOTE Based on the type of cluster and number of hosts, vSAN creates and assigns a default datastore policy for best capacity utilization after the cluster configuration is completed. Policy details can be viewed in the vSphere Client (Policies and Profiles › VM Storage Policies).</p> <p>Click Next.</p>

9. If you selected NFS storage for the vSphere cluster, the NFS Storage page appears.
 - a) Enter a name for the NFS datastore.
 - b) Enter the path to the NFS share.
 - c) Enter the IP address of the NFS server.
 - d) Click **Next**.
10. If you selected VMFS on FC storage for the vSphere cluster, enter the name of the VMFS on FC datastore and click **Next**.
11. If you selected vVols storage for the vSphere cluster, the vVols storage page appears.
 - a) Select a VASA protocol type.
vVols supports FC, NFS, and iSCSI storage protocol types.
 - b) Select a VASA provider name.
 - c) Select a storage container.
 - d) Select a VASA user.

- e) Enter a datastore name.
 - f) Click **Next**.
12. On the Host Selection page, select hosts for the vSphere cluster.

You can use the toggle button to turn **Skip failed hosts during cluster creation** off or on. When this option is off, cluster creation will fail if you select an unhealthy host. When this option is on, cluster creation will succeed if you selected enough healthy hosts to meet the minimum requirements for a new cluster.

To view DPU-backed hosts, activate the **Network Offloading** toggle. Do not activate the toggle if you want to select hosts that are not DPU-backed.

NOTE

The toggle is only available if the workload domains uses vLCM images and there are unassigned DPU-backed hosts available.

If you are using DPU-backed hosts, select the **DPU Vendor**.

NOTE

All hosts in a cluster must use DPUs from the same vendor.

When you use the SDDC Manager UI to add a cluster, all hosts must be associated with the same network pool. The VMware Cloud Foundation API supports adding hosts from different network pools to workload domain clusters, as long as those network pools have the same VLAN ID and MTU settings.

The hosts must be commissioned with the same storage type that you selected for the cluster. For example, select hosts commissioned for vSAN Max storage for a vSAN Max cluster.

When you have selected the minimum number of hosts required for this cluster, the **Next** button is enabled.

13. Click **Next**.
14. On the Switch Configuration page, select either a preconfigured switch profile to apply to the cluster or select the option to create a custom switch configuration.
- a) When creating a custom switch configuration, specify the Distributed Switch Name, MTU, number of VDS uplinks, and uplink mapping.
 - b) Click **Configure Network Traffic** and select the network traffic type to configure.

You must configure all required network traffic types.

Network Traffic Type	Configuration
Management, vMotion, vSAN, Public	Enter the Distributed PortGroup Name, select the Load Balancing policy, and configure the uplinks. Click Save Configuration .
NSX	Select the Operational Mode (Standard, Enhanced Datapath - Standard, Enhanced Datapath - Performance) and Transport Zone Type(s) (NSX-Overlay, NSX-VLAN). NOTE For a VI workload domain with DPU-backed hosts, you must select Enhanced Datapath - Standard .

Table continued on next page

Continued from previous page

Network Traffic Type	Configuration
	<p>For NSX-Overlay:</p> <ul style="list-style-type: none"> • Enter the NSX-Overlay Transport Zone Name. • Enter the VLAN ID and select the IP Allocation (DHCP or Static IP Pool). <p>For NSX-VLAN:</p> <ul style="list-style-type: none"> • Enter the NSX-VLAN Transport Zone Name. <p>Configure the:</p> <ul style="list-style-type: none"> • NSX transport node settings • Uplink mapping settings • NSX uplink profile settings <p>Click Save Configuration.</p>

c) Click **Create Distributed Switch**.

NOTE

You cannot proceed until all mandatory traffic types are configured.

d) Click **Next**.

15. On the Licenses page, choose a licensing option.

Option	Description
License Now	Select the vSphere and vSAN (if applicable) license to apply to this cluster.
License Later	<p>vSphere and vSAN (if applicable) components are deployed in evaluation mode.</p> <p>IMPORTANT</p> <p>After the cluster is created, you must switch to licensed mode by:</p> <ul style="list-style-type: none"> • Adding component license keys in the SDDC Manager UI. See Add a Component License Key in the SDDC Manager UI. Or, • Adding a solution license key in the vSphere Client. See Managing vSphere Licenses. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See Configure License Settings for a vSAN Cluster.

16. Click **Next**.


17. On the Review page, review the vSphere cluster details and click **Finish**.

NOTE

Multiple VMkernels are created to test the vMotion network, which may cause changes in the MAC addresses and IP address relations. If MAC address filtering is enabled on your physical infrastructure, this may cause issues such as vMotion network connectivity validation failure.

The details page for the workload domain appears with the following message: *Adding a new cluster is in progress.* When this process completes, the vSphere cluster appears in the Clusters tab in the details page for the workload domain.

If **Skip failed hosts during cluster creation** is on and one or more of your selected hosts is unhealthy, the cluster creation task still succeeds as long as you have at least the minimum number of healthy hosts. The Tasks panel reports that a host was skipped and includes a link with more details.



If, after skipping failed hosts, there are not enough healthy hosts to create a cluster, the task fails.

To add another cluster in parallel, click **Add Cluster** again and repeat the above steps.

Shrink a Workload Domain

You can reduce the management domain or a VI workload domain by removing a host from a vSphere cluster in the workload domain or by deleting a vSphere cluster.

Remove a Host from a vSphere Cluster in a Workload Domain

You can remove a host from a vSphere cluster in the management domain or a VI workload domain through the Workload Domains page in SDDC Manager UI.

- Delete any workload virtual machines created outside VMware Cloud Foundation or move them to another host.
- You cannot remove a host from a vSphere cluster if that host is hosting an NSX Edge node.
 - Move the NSX Edge node to another host in the same vSphere cluster. You cannot move an NSX Edge node to a host that already hosts another NSX Edge node.
 - If you cannot move the NSX Edge node, delete the NSX Edge node on the host. You cannot delete NSX Edge nodes if doing so would result in an NSX Edge cluster with fewer than two NSX Edge nodes.
- You cannot remove a host from a vSphere cluster if that vSphere cluster was selected for NSX Edge node placement and the NSX Edge node deployment is still in pending state.

Before you remove a host from a vSphere cluster, ensure that you have enough hosts remaining to facilitate the configured vSAN availability. Failure to do so might result in the datastore being marked as read-only or in data loss.

NOTE

You can add hosts to or remove hosts from multiple different vSphere clusters in parallel. For example, you remove three hosts from Cluster A, and while that task is running, you can start a separate task to remove (or add) four hosts to Cluster B. See [VMware Configuration Maximums](#) for information about the maximum number of add/remove hosts tasks that you can run in parallel.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the workload domains table, click the name of the workload domain that you want to modify.
3. Click the **Clusters** tab.
4. Click the name of the cluster from which you want to remove a host.
5. Click the **Hosts** tab.
6. Select the host to remove and click **Remove Selected Hosts**.

An alert appears, asking you to confirm or cancel the action. If the removal results in the number of hosts in the vSphere cluster being less than the minimum number of required hosts, you must click **Force Remove** to remove the host.

7. Click **Remove** to confirm the action.

The details page for the cluster appears with a message indicating that the host is being removed. When the removal process is complete, the host is removed from the hosts table.

Decommission the host. See [Decommission Hosts](#). After you decommission a host, you must re-image it and commission it before you can use it again.

Delete a vSphere Cluster from a Workload Domain

You can delete a vSphere cluster from the management domain or from a VI workload domain. vSAN datastores on the ESXi hosts in the deleted cluster are destroyed, while other types of storage are unmounted.

- If remote vSAN datastores are mounted on the cluster, the cluster cannot be deleted. To delete such clusters, you must first migrate any VMs from the remote datastore to the local datastore and then unmount the vSAN remote datastores from vCenter Server.
- Migrate or back up the VMs and data on the datastore associated with the cluster to another location if you want the ability to restore them later.
- Delete the NSX Edge clusters hosted on the vSphere cluster or shrink the NSX Edge cluster by deleting NSX Edge nodes hosted on the vSphere cluster. You cannot delete NSX Edge nodes if doing so would result in an NSX Edge cluster with fewer than two NSX Edge nodes. For information about deleting an NSX Edge cluster, see [KB 78635](#). For information about removing an NSX Edge node from an NSX Edge cluster, see [Remove Edge Nodes from an NSX Edge Cluster](#).

You cannot delete the last cluster in a workload domain. Instead, delete the workload domain. See [Delete a VI Workload Domain](#).

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. Click the name of the workload domain that contains the vSphere cluster you want to delete.
3. Click the **Clusters** tab.
4. Click the vertical ellipsis (three dots) next to the cluster name and click **Delete Cluster**.

Add Host

Mount Remote Datastore

Delete Cluster

Rename Cluster

5. Click **Delete Cluster** to confirm that you want to delete the cluster.

You cannot perform additional workload domain tasks until the Delete Cluster workflow has completed.

Rename a Workload Domain

You can rename any workload domain from within the SDDC Manager UI.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. Click the vertical ellipsis (three dots) in the Domain row for the workload domain you want to rename and click **Rename Domain**.
3. Enter a new name for the workload domain and click **Rename**.

vSphere Cluster Management

You can view vSphere cluster details from the SDDC Manager UI and rename the vSphere Cluster using the vSphere Client if required.

View vSphere Cluster Details

The cluster summary page displays high level information about the vSphere cluster as well as the hosts that form that cluster. CPU, memory, and storage utilization are also displayed.

1. In the navigation pane, click **Inventory** > **Workload Domain**.
2. In the workload domains table, click the name of a workload domain.
3. Click the **Clusters** tab.
4. In the clusters table, click the name of a vSphere cluster.

The screenshot shows the vSphere Cluster Management UI for a cluster named 'vi-cluster1'. The page is divided into several sections:

- Resource Usage:** Displays CPU usage at 8.94 GHz (used of 50.28 GHz, 41.34 GHz free) and Memory usage at 56.12 GB (used of 351.55 GB, 295.43 GB free). A dropdown menu is set to 'NUM' and '%'. There is an 'ASSIGN' button.
- Tags:** Shows 'No tags assigned' and an 'ASSIGN' button.
- Storage:** Shows 'Local vSAN' with a progress bar indicating 'Used 36.43 GB / 858.23 GB (Free 821.8 GB)'. Below this, it shows 'Datastore Name: vi-cluster1-vSanDatastore' and 'FTT: 0'.

The cluster detail page appears. The tabs on the page display additional information as described in the table below.

Tab	Information Displayed
Summary	Displays information about resource usage, storage, and cluster tags.
Hosts	Details about the ESXi hosts in the vSphere cluster. You can click a name in the FQDN column to access the host summary page.

You can add or remove a host, or access the vSphere Client from this page.

Rename a Cluster in the SDDC Manager UI

You can rename a cluster managed by SDDC Manager in a Management Workload Domain. The SDDC Manager UI is updated with the new name.

Ensure that you do not rename a cluster in the following conditions:

- When the cluster belongs to a workflow that is in progress.
- When the cluster belongs to a failed VI workload domain workflow, cluster workflow or host workflow. If you try to rename a cluster that belongs to a failed workflow, restart of the failed workflow will not be supported.

1. On the SDDC Manager Dashboard, click **Inventory › Workload Domains**.
2. Click a workload domain.
3. Under the **Clusters** tab, click a cluster that you want to rename.
4. On the right side of the cluster's name, click **ACTIONS › Rename Cluster**.

You can also click the vertical ellipsis (three dots) in the clusters table for the cluster you want to rename and click **Rename Cluster**.

The **Rename Cluster** window appears.

5. In the **New Cluster Name** textbox, enter a new name for the cluster and click **RENAME**.
6. Click **DONE**.

In the **Tasks** panel, you can see the description and track the status of your newly renamed cluster.

Mount a Remote vSAN Datastore

Remote datastore sharing enables vSAN clusters to share their datastores with other clusters. Mounting a remote vSAN datastore is a cluster-wide configuration. When you mount a remote vSAN datastore to a vSAN cluster, it is available to all hosts in the cluster. Use the SDDC Manager UI to remotely mount datastores from other vSAN clusters.

At least one vSAN datastore of a supported type must exist in the workload domain.

Cluster Type	Supported Remote Datastores
vSAN OSA	vSAN OSA
vSAN ESA	vSAN ESA, vSAN Max
vSAN Compute Cluster	vSAN OSA, vSAN ESA, vSAN Max

A vSAN cluster can mount a remote vSAN datastore from another vSAN cluster in the same workload domain. Once mounted, the cluster consumes storage resources from the remote datastore. Remote clusters allows you to share storage resources from underutilized clusters or from vSAN Max clusters, which are designed specifically to provide storage resources.

vSAN compute clusters provide compute resources only and require a remote datastore for storage resources.

NOTE

- Two or more clusters can share the same remote vSAN datastore.
- A cluster can mount more than one remote vSAN datastore, as long as all the datastores are of the same type.

1. In the navigation pane, click **Inventory › Workload Domains**.
2. In the domain table, click the name of a domain.
3. Click the **Clusters** tab.
4. In the clusters table, click the cluster name.
5. Click the **Actions** dropdown to the right of the cluster name and select **Mount Remote Datastore**.
6. Select a remote datastore from the list of available vSAN clusters and click **Next**.
7. Review the details and click **Finish**.

The remote datastore shows as mounted on the existing vSAN cluster.

Unmount a Remote vSAN Datastore

Use the SDDC Manager UI to unmount a remote vSAN datastore.

To unmount a remote vSAN datastore, make sure no virtual machines, including the system vCLS virtual machines, are residing on it.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the domain table, click the name of a domain.
3. Click the **Clusters** tab.
4. In the clusters table, click the cluster name.
5. Click the **Unmount**.
6. Select the remote vSAN datastore to unmount from the cluster and click **Next**.
7. Review the details and click **Finish**.

The remote vSAN datastore is unmounted from the cluster.

Tag Management

A tag is a label that you can apply to objects in the vSphere inventory. You can use tags to capture a variety of metadata about your vSphere inventory and to organize and retrieve objects quickly. You create tags and categories in the vSphere Client and then assign or remove tags for your workload domains, clusters, and hosts in the SDDC Manager UI.

See [vSphere Tags](#) for more information about how to create and manage tags and categories.

If multiple vCenter Server instances in your VMware Cloud Foundation deployment are configured to use Enhanced Linked Mode, tags and tag categories are replicated across all these vCenter Server instances. This is the case for all VI workload domains that are joined to the same SSO domain as the management domain. Isolated VI workload domains, that do not share the management SSO domain, do not share its tags and categories.

Tag a Workload Domain

You can assign a tag to your Workload Domain from the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** and click the Workload Domain.
2. Under the **Summary** › **Tags** tile window, click **ASSIGN**.

Assign Tag

[OPEN VSPHERE TAG MANAGEMENT](#) ✕

Assign tags to the workload domain "nsxt-vi". To create new tags or new tag categories, go to vSphere Tag Management.

<input type="checkbox"/>	Tag name ↑ ▼	Tag category ▼	Tag associations per category ▼
<input type="checkbox"/>	workload-domain-production	Workload Domains	One
<input type="checkbox"/>	workload-domain-test	Workload Domains	One
			1 - 2 of 2 Tags

CANCEL

ASSIGN

3. Select tags and click **ASSIGN**.

NOTE

If there are no tags shown in the **Assign Tag** window, click **OPEN VSPHERE TAG MANAGEMENT** that redirects you to vSphere Client, to create new tags and tag categories. See [vSphere Tags](#) for more information on the tagging functionality.

Remove a Tag from your Workload Domain

You can remove a tag from your workload domain in the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** and click the workload domain.
2. Under the **Summary** › **Tags** tile window, you will see tags listed with a cross mark beside the tag names.
3. Click the cross mark of a tag that you want to remove in the **Tags** tile window.
4. The **Remove Tag** window appears. Click **REMOVE**.

Tag a Cluster


You can assign a tag to your cluster from the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** › **Workload Domain** › **Clusters** tab and click on the cluster.
2. Under the **Summary** › **Tags** tile window, click **ASSIGN**.

Assign Tag

[OPEN VSPHERE TAG MANAGEMENT](#) ✕

Assign tags to the cluster "vi-cluster1". To create new tags or new tag categories, go to vSphere Tag Management.

<input type="checkbox"/>	Tag name ↑ ▼	Tag category ▼	Tag associations per category ▼
<input type="checkbox"/>	dev-cluster	Cluster Users	One
<input type="checkbox"/>	marketing-cluster	Cluster Users	One
			1 - 2 of 2 Tags

CANCEL

ASSIGN

3. Select tags and click **ASSIGN**.

NOTE

If there are no tags shown in the **Assign Tag** window, click **OPEN VSPHERE TAG MANAGEMENT** that redirects you to vSphere Client, to create new tags and tag categories. See [vSphere Tags](#) for more information on the tagging functionality.

Remove a Tag from your Cluster

You can remove a tag from your cluster in the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** › **Workload Domain** › **Clusters** tab and click on the cluster.
2. Under the **Summary** › **Tags** tile window, you will see tags listed with a cross mark beside the tag names.
3. Click the cross mark of a tag that you want to remove in the **Tags** tile window.
4. The **Remove Tag** window appears. Click **REMOVE**.

Tag a Host

You can assign a tag to your host from the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Hosts** tab and click on the host.
2. Under the **Summary** › **Tags** tile window, click **ASSIGN**.

Assign Tag

[OPEN VSPHERE TAG MANAGEMENT](#) ✕

Assign tags to the host "esxi-1.vrack.vsphere.local". To create new tags or new tag categories, go to vSphere Tag Management.

<input type="checkbox"/>	Tag name	↑ ▼	Tag category	▼	Tag associations per category	▼
<input type="checkbox"/>	mgmt-hosts		Workload Domain Hosts		One	
<input type="checkbox"/>	vi-wld-hosts		Workload Domain Hosts		One	
						1 - 2 of 2 Tags

CANCEL

ASSIGN

3. Select tags and click **ASSIGN**.

NOTE

If there are no tags shown in the **Assign Tag** window, click **OPEN VSPHERE TAG MANAGEMENT** that redirects you to vSphere Client, to create new tags and tag categories. See [vSphere Tags](#) for more information on the tagging functionality.

Remove a Tag from your Host

You can remove a tag from your host in the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Hosts** tab and click on the host.
2. Under the **Summary** › **Tags** tile window, you will see tags listed with a cross mark beside the tag names.
3. Click the cross mark of a tag that you want to remove in the **Tags** tile window.
4. The **Remove Tag** window appears. Click **REMOVE**.

Manage Workload Domain Configuration Drift Between vCenter Server and SDDC Manager

If you make changes to your VMware Cloud Foundation configuration using the vSphere client, you can use the sync workflow in the VCF Import Tool to manually update the SDDC Manager with any out-of-band changes applied from vCenter Server.

You must have the latest version of the VCF Import Tool on the SDDC Manager appliance.

If the SDDC Manager and vCenter Server inventories get out of sync, some SDDC Manager workflows may be blocked. You can use this procedure to unblock these tasks.

You can run the sync workflow on vSphere environments that you imported/converted to SDDC Manager workload domains, as well as on workload domains deployed from SDDC Manager.

The sync workflow requires you to specify a workload domain to sync. If you need to sync multiple workload domains, you will have to run the sync workflow multiple times.

1. If the VCF Import Tool is already deployed on the SDDC Manager appliance, skip to step 3. Otherwise, download it now.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads** › **VMware Cloud Foundation** .
 - b) Click the version 5.2.

- c) Click **Drivers & Tools**.
 - d) Click the download icon for the VCF Import Tool.
 - e) Copy `vcf-brownfield-import-<version>.tar.gz` to the SDDC Manager appliance.
2. Copy `vcf-brownfield-import-<version>.tar.gz` to the SDDC Manager appliance and extract the archive.
 - a) SSH in to the SDDC Manager appliance using the `vcf` user account.
 - b) Create the `vcfimport` directory.


```
mkdir /home/vcf/vcfimport
```
 - c) Copy `vcf-brownfield-import-<version>.tar.gz` to the `/home/vcf/vcfimport` directory.
 - d) Extract the contents of `vcf-brownfield-import-<version>.tar.gz`.


```
tar -xvf vcf-brownfield-import-<version>.tar.gz
```
 3. Run the `vcf_brownfield.py` script with the `sync` parameter.
 - a) SSH in to the SDDC Manager appliance using the `vcf` user account.
 - b) Navigate to the location of the `vcf_brownfield.py` script and run the following command:


```
python3 vcf_brownfield.py sync --domain-name <selected-domain-name>
```

 Replace `<selected-domain-name>` with the name of the workload domain you want to sync.

Managing NSX Edge Clusters in VMware Cloud Foundation

An NSX Edge cluster with 2-tier routing provides north-south routing and network services in the management domain and VI workload domains. Add multiple NSX Edge clusters to a workload domain for scalability and resiliency.

An NSX Edge cluster is a logical grouping of NSX Edge nodes run on a vSphere cluster. NSX supports a 2-tier routing model.

Component	Connectivity	Description
Tier-0 logical router	Northbound	The tier-0 logical router connects to one or more physical routers or layer 3 switches and serves as a gateway to the physical infrastructure.
	Southbound	The tier-0 logical router connects to one or more tier-1 logical routers or directly to one or more logical switches.
Tier-1 logical router	Northbound	The tier-1 logical router connects to a tier-0 logical router.
	Southbound	The tier-1 logical router connects to one or more logical switches.

By default, workload domains do not include any NSX Edge clusters and workloads are isolated, unless VLAN-backed networks are configured in vCenter Server. Add one or more NSX Edge clusters to a workload domain to provide software-defined routing and network services.

NOTE

You must create an NSX Edge cluster on the default management vSphere cluster in order to deploy VMware Aria Suite products.

You can add multiple NSX Edge clusters to the management or the VI workload domains for scalability and resiliency. For VMware Cloud Foundation configuration maximums refer to the [VMware Configuration Maximums](#) website.

NOTE

Unless explicitly stated in this matrix, VMware Cloud Foundation supports the configuration maximums of the underlying products. Refer to the individual product configuration maximums as appropriate.

The north-south routing and network services provided by an NSX Edge cluster created for a workload domain are shared with all other workload domains that use the same NSX Manager cluster.

Prerequisites for an NSX Edge Cluster

Before you deploy an NSX Edge cluster on a workload domain, review the prerequisites.

- The workload domain must have NSX deployed.
- Verify that separate VLANs and subnets are available for the NSX host overlay VLAN and NSX Edge overlay VLAN. You cannot use DHCP for the NSX Edge overlay VLAN.
- Verify that the NSX host overlay VLAN and NSX Edge overlay VLAN are routed to each other.
- For dynamic routing, set up two Border Gateway Protocol (BGP) peers on Top of Rack (ToR) switches with an interface IP, BGP autonomous system number (ASN), and BGP password.
- Reserve a BGP ASN to use for the NSX Edge cluster's Tier-0 gateway.
- Verify that DNS entries for the NSX Edge nodes are populated in the customer-managed DNS server.
- The vSphere cluster hosting an NSX Edge cluster must include hosts with identical management, uplink, NSX Edge overlay TEP, and NSX Edge overlay TEP networks (L2 uniform).
- If the vSphere cluster hosting the NSX Edge nodes has hosts with a DPU device:
 - Enable SR-IOV in the BIOS and in the vSphere Client (if required by your DPU vendor).
 - Ensure that UPT is enabled for the DPU-backed NICs.

IMPORTANT

You cannot use SDDC Manager to deploy an NSX Edge cluster on a vSphere cluster with dual DPU hosts configured for high-availability. Use NSX Manager instead.

- The management network and management network gateway for the NSX Edge nodes must be reachable from the NSX host overlay and NSX Edge overlay VLANs.

NOTE

VMware Cloud Foundation 4.5 and later support deploying an NSX Edge cluster on a vSphere cluster that is stretched. Edge nodes are placed on ESXi hosts in the first availability zone (AZ1) during NSX Edge cluster deployment.

Deploy an NSX Edge Cluster

Deploy an NSX Edge cluster to provide north-south routing and network services to a workload domain.

See [Prerequisites for an NSX Edge Cluster](#).

SDDC Manager does not enforce rack failure resiliency for NSX Edge clusters. Make sure that the number of NSX Edge nodes that you add to an NSX Edge cluster, and the vSphere clusters to which you deploy the NSX Edge nodes, are sufficient to provide NSX Edge routing services in case of rack failure.

After you create an NSX Edge cluster, you can use SDDC Manager to expand or shrink it by adding or deleting NSX Edge nodes.

This procedure describes how to use SDDC Manager to create an NSX Edge cluster with NSX Edge node virtual appliances. If you have latency intensive applications in your environment, you can deploy NSX Edge nodes on bare-metal servers. See [Deployment of VMware NSX-T Edge Nodes on Bare-Metal Hardware for VMware Cloud Foundation 4.0.x](#).

NOTE

If you deploy the NSX Edge cluster with the incorrect settings or need to delete an NSX Edge cluster for another reason, see [KB 78635](#).

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Select **Actions** > **Add Edge Cluster**.
4. Verify the prerequisites, select **Select All**, and click **Begin**.
5. Enter the configuration settings for the NSX Edge cluster and click **Next**.

Setting	Description
Edge Cluster Name	Enter a name for the NSX Edge cluster.
MTU	Enter the MTU for the NSX Edge cluster. The MTU can be 1600-9000.
Tier-0 Router Name	Enter a name for the tier-0 gateway.
Tier-1 Router Name	Enter a name for the tier-1 gateway.
Edge Cluster Profile Type	Select Default or, if your environment requires specific Bidirectional Forwarding Detection (BFD) configuration, select Custom .
Edge Cluster Profile Name	Enter an NSX Edge cluster profile name. (Custom Edge cluster profile only)
BFD Allowed Hop	Enter the number of multi-hop Bidirectional Forwarding Detection (BFD) sessions allowed for the profile. (Custom Edge cluster profile only)
BFD Declare Dead Multiple	Enter the number of number of times the BFD packet is not received before the session is flagged as down. (Custom Edge cluster profile only)
BFD Probe Interval (milliseconds)	BFD is detection protocol used to identify the forwarding path failures. Enter a number to set the interval timing for BFD to detect a forwarding path failure. (Custom Edge cluster profile only)
Standby Relocation Threshold (minutes)	Enter a standby relocation threshold in minutes. (Custom Edge cluster profile only)
Edge Root Password	Enter and confirm the password to be assigned to the root account of the NSX Edge appliance.
Edge Admin Password	Enter and confirm the password to be assigned to the admin account of the NSX Edge appliance.
Edge Audit Password	Enter and confirm the password to be assigned to the audit account of the NSX Edge appliance.

NSX Edge cluster passwords must meet the following requirements:

- At least 12 characters
- At least one lower-case letter
- At least one upper-case letter
- At least one digit
- At least one special character (!, @, ^, =, *, +)
- At least five different characters
- No dictionary words
- No palindromes
- More than four monotonic character sequence is not allowed

6. Specify the use case details and click **Next**.

Setting	Description
Use Case	<ul style="list-style-type: none"> Select Kubernetes - Workload Management to create an NSX Edge cluster that complies with the requirements for deploying vSphere IaaS Control Plane. See VMware Cloud Foundation with VMware Tanzu . If you select this option, you cannot modify the NSX Edge form factor or Tier-0 service high availability settings. Select Application Virtual Networks to create an NSX Edge cluster that complies with the requirements for deploying VMware Aria Suite components. See Deploying Application Virtual Networks in VMware Cloud Foundation. <p style="text-align: center;">NOTE Management domain only.</p> <ul style="list-style-type: none"> Select Custom if you want an NSX Edge cluster with a specific form factor or Tier-0 service high availability setting.
Edge Form Factor	<ul style="list-style-type: none"> Small: 4 GB memory, 2 vCPU, 200 GB disk space. The NSX Edge Small VM appliance size is suitable for lab and proof-of-concept deployments. Medium: 8 GB memory, 4 vCPU, 200 GB disk space. The NSX Edge Medium appliance size is suitable for production environments with load balancing. Large: 32 GB memory, 8 vCPU, 200 GB disk space. The NSX Edge Large appliance size is suitable for production environments with load balancing. XLarge: 64 GB memory, 16 vCPU, 200 GB disk space. The NSX Edge Extra Large appliance size is suitable for production environments with load balancing.
Tier-0 Service High Availability	<p>In the active-active mode, traffic is load balanced across all members. In active-standby mode, all traffic is processed by an elected active member. If the active member fails, another member is elected to be active. Workload Management requires Active-Active.</p> <p>Some services are only supported in Active-Standby: NAT, load balancing, stateful firewall, and VPN. If you select Active-Standby, use exactly two NSX Edge nodes in the NSX Edge cluster.</p>
Tier-0 Routing Type	<p>Select Static or EBGP to determine the route distribution mechanism for the tier-0 gateway. If you select Static, you must manually configure the required static routes in NSX Manager. If you select EBGP, VMware Cloud Foundation configures eBGP settings to allow dynamic route distribution.</p>
ASN	<p>Enter an autonomous system number (ASN) for the NSX Edge cluster. (for EBGP only)</p>

7. Enter the configuration settings for the first NSX Edge node and click **Add Edge Node**.

Setting	Description
Edge Node Name (FQDN)	Enter the FQDN for the NSX Edge node. Each node must have a unique FQDN.
Cluster	<p>Select a vSphere cluster to host the NSX Edge node. You can select a standard vSphere cluster or a stretched vSphere cluster, but all the NSX Edge nodes in an NSX Edge cluster must be hosted on vSphere clusters of the same type.</p> <p>NOTE If the vSphere cluster you select already hosts management virtual machines that are connected to the host Management port group, the VM Management Portgroup VLAN and VM Management Portgroup VLAN settings are not available.</p>
Cluster Type	<p>Select L2 Uniform if all hosts in the vSphere cluster have identical management, uplink, host TEP, and Edge TEP networks. Select L2 non-uniform and L3 if any of the hosts in the vSphere cluster have different networks.</p> <p>IMPORTANT VMware Cloud Foundation does not support Edge cluster creation on L2 non-uniform and L3 vSphere clusters.</p>
First NSX VDS Uplink	<p>Click Advanced Cluster Settings to map the first NSX Edge node uplink network interface to a physical NIC on the host, by specifying the ESXi uplink. The default is <code>uplink1</code>.</p> <p>When you create an NSX Edge cluster, SDDC Manager creates two trunked VLAN port groups. The information you enter here determines the active uplink on the first VLAN port group. If you enter <code>uplink3</code>, then <code>uplink3</code> is the active uplink and the uplink you specify for the second NSX VDS uplink is the standby uplink.</p> <p>The uplink must be prepared for overlay use.</p>
Second NSX VDS Uplink	<p>Click Advanced Cluster Settings to map the second NSX Edge node uplink network interface to a physical NIC on the host, by specifying the ESXi uplink. The default is <code>uplink2</code>.</p> <p>When you create an NSX Edge cluster, SDDC Manager creates two trunked VLAN port groups. The information you enter here determines the active uplink on the second VLAN port group. If you enter <code>uplink4</code>, then <code>uplink4</code> is the active uplink and the uplink you specify for the first NSX VDS uplink is the standby uplink.</p> <p>The uplink must be prepared for overlay use.</p>

Table continued on next page

Continued from previous page

Setting	Description
Management IP (CIDR)	Enter the management IP for the NSX Edge node in CIDR format. Each node must have a unique management IP.
Management Gateway	Enter the IP address for the management network gateway.
VM Management Portgroup VLAN	If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group VLAN is displayed and cannot be edited. If the VM Management port group does not exist on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, enter a VLAN ID to create a new VM Management port group or click Use ESXi Management VMK's VLAN to use the host Management Network VLAN to create a new VM Management port group.
VM Management Portgroup Name	If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group name is displayed and cannot be edited. Otherwise, type a name for the new port group.
Edge TEP 1 IP (CIDR)	Enter the CIDR for the first NSX Edge TEP. Each node must have a unique Edge TEP 1 IP. NOTE It is possible to configure Edge TEPs using an NSX IP pool instead of static addresses. IP pools may only be specified when using the VCF API only, not the UI.
Edge TEP 2 IP (CIDR)	Enter the CIDR for the second NSX Edge TEP. Each node must have a unique Edge TEP 2 IP. The Edge TEP 2 IP must be different than the Edge TEP 1 IP.
Edge TEP Gateway	Enter the IP address for the NSX Edge TEP gateway.
Edge TEP VLAN	Enter the NSX Edge TEP VLAN ID.
First Tier-0 Uplink VLAN	Enter the VLAN ID for the first uplink. This is a link from the NSX Edge node to the first uplink network.
First Tier-0 Uplink Interface IP (CIDR)	Enter the CIDR for the first uplink. Each node must have unique uplink interface IPs.
Peer IP (CIDR)	Enter the CIDR for the first uplink peer. (EBGP only)
Peer ASN	Enter the ASN for the first uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only).
Second Tier-0 Uplink VLAN	Enter the VLAN ID for the second uplink. This is a link from the NSX Edge node to the second uplink network.

Table continued on next page

Continued from previous page

Setting	Description
Second Tier-0 Uplink Interface IP (CIDR)	Enter the CIDR for the second uplink. Each node must have unique uplink interface IPs. The second uplink interface IP must be different than the first uplink interface IP.
Peer IP (CIDR)	Enter the CIDR for the second uplink peer. (EBGP only)
ASN Peer	Enter the ASN for the second uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only).

8. Click **Add More Edge Nodes** to enter configuration settings for additional NSX Edge nodes.

A minimum of two NSX Edge nodes is required. NSX Edge cluster creation allows up to 8 NSX Edge nodes if the Tier-0 Service High Availability is Active-Active and two NSX Edge nodes per NSX Edge cluster if the Tier-0 Service High Availability is Active-Standby.


NOTE

All Edge nodes in the NSX Edge cluster must use the same VM Management port group VLAN and name.

9. When you are done adding NSX Edge nodes, click **Next**.
10. Review the summary and click **Next**.
SDDC Manager validates the NSX Edge node configuration details.
11. If validation fails, use the **Back** button to edit your settings and try again.
To edit or delete any of the NSX Edge nodes, click the three vertical dots next to an NSX Edge node in the table and select an option from the menu.
12. If validation succeeds, click **Finish** to create the NSX Edge cluster.
You can monitor progress in the Tasks panel.

The following example shows a scenario with sample data. You can use the example to guide you in creating NSX Edge clusters in your environment. Refer to the [Planning and Preparation Workbook](#) for a complete list of sample values for creating an NSX Edge cluster.

Figure 22: Two-node NSX Edge cluster in a single rack



In NSX Manager, you can create segments connected to the NSX Edge cluster's tier-1 gateway. You can connect workload virtual machines to these segments to provide north-south and east-west connectivity.

Add Edge Nodes to an NSX Edge Cluster

You can add NSX Edge nodes to an NSX Edge Cluster that you created with SDDC Manager.

- Verify that separate VLANs and subnets are available for the NSX host overlay VLAN and NSX Edge overlay VLAN. You cannot use DHCP for the NSX Edge overlay VLAN.
- Verify that the NSX host overlay VLAN and NSX Edge overlay VLAN are routed to each other.

- For dynamic routing, set up two Border Gateway Protocol (BGP) peers on Top of Rack (ToR) switches with an interface IP, BGP autonomous system number (ASN), and BGP password.
- Reserve a BGP ASN to use for the NSX Edge cluster's Tier-0 gateway.
- Verify that DNS entries for the NSX Edge nodes are populated in the customer-managed DNS server.
- The vSphere cluster hosting the NSX Edge nodes must include hosts with identical management, uplink, NSX Edge overlay TEP, and NSX Edge overlay TEP networks (L2 uniform).
- The vSphere cluster hosting the NSX Edge nodes must have the same pNIC speed for NSX-enabled VDS uplinks chosen for Edge overlay.
- All NSX Edge nodes in an NSX Edge cluster must use the same set of NSX-enabled VDS uplinks. These uplinks must be prepared for overlay use.
- The NSX Edge cluster must be **Active**.
- The NSX Edge cluster must be hosted on one or more vSphere clusters from the same workload domain.
- If the vSphere cluster hosting the NSX Edge nodes has hosts with a DPU device:
 - Enable SR-IOV in the BIOS and in the vSphere Client (if required by your DPU vendor).
 - Ensure that UPT is enabled for the DPU-backed NICs.

You might want to add NSX Edge nodes to an NSX Edge cluster, for:

- Rack failure resiliency
- When the Tier-0 Service High Availability is Active-Standby and you require more than two NSX Edge nodes for services.

NOTE

Only two of the NSX Edge nodes can have uplink interfaces, but you can add more nodes without uplink interfaces.

- When the Tier-0 Service High Availability is Active-Active and you require more than 8 NSX Edge nodes for services.
- When you add Supervisor Clusters to a Workload Management workload domain and need to support additional tier-1 gateways and services.

The available configuration settings for a new NSX Edge node vary based on:

- The Tier-0 Service High Availability setting (Active-Active or Active-Standby) of the NSX Edge cluster.
- The Tier-0 Routing Type setting (static or EBGp) of the NSX Edge cluster.
- Whether the new NSX Edge node is going to be hosted on the same vSphere cluster as the existing NSX Edge nodes (in-cluster) or on a different vSphere cluster (cross-cluster).

1. In the navigation pane, click **Inventory > Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Click the **Edge Clusters** tab.
4. Click the vertical ellipsis menu for the Edge Cluster you want to expand and select **Expand Edge Cluster**.
5. Verify the prerequisites, select **Select All**, and click **Begin**.
6. Enter and confirm the passwords for the NSX Edge cluster.
7. Enter a name to create a new tier-1 gateway.
8. Enter the configuration settings for the new NSX Edge node and click **Add Edge Node**.

Setting	Description
Edge Node Name (FQDN)	Enter the FQDN for the NSX Edge node. Each node must have a unique FQDN.
Cluster	Select a vSphere cluster to host the NSX Edge node. If the workload domain has multiple vSphere clusters, you can select the vSphere cluster hosting the existing NSX Edge nodes (in-cluster expansion) or select a

Table continued on next page

Continued from previous page

Setting	Description
	<p>different vSphere cluster to host the new NSX Edge nodes (cross-cluster expansion).</p> <p>NOTE If the vSphere cluster you select already hosts management virtual machines that are connected to the host Management port group, the VM Management Portgroup VLAN and VM Management Portgroup VLAN settings are not available.</p>
Cluster Type	<p>Select L2 Uniform if all hosts in the vSphere cluster have identical management, uplink, host TEP, and Edge TEP networks.</p> <p>Select L2 non-uniform and L3 if any of the hosts in the vSphere cluster have different networks.</p> <p>IMPORTANT VMware Cloud Foundation does not support Edge cluster creation on L2 non-uniform and L3 vSphere clusters.</p>
Management IP (CIDR)	Enter the management IP for the NSX Edge node in CIDR format. Each node must have a unique management IP.
Management Gateway	Enter the IP address for the management network gateway.
VM Management Portgroup VLAN	<p>For in-cluster expansion, the new Edge node uses the same VM Management port group VLAN as the other Edge nodes in the Edge cluster.</p> <p>For cross-cluster expansion:</p> <ul style="list-style-type: none"> • If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group VLAN is displayed and cannot be edited. • If the VM Management port group does not exist on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, enter a VLAN ID to create a new VM Management port group or click Use ESXi Management VMK's VLAN to use the host Management Network VLAN for the VM Management port group.
VM Management Portgroup Name	<p>For in-cluster expansion, the new Edge node uses the same VM Management port group name as the other Edge nodes in the Edge cluster.</p> <p>For cross-cluster expansion:</p> <ul style="list-style-type: none"> • If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the

Table continued on next page

Continued from previous page

Setting	Description
	<p>VM Management port group name is displayed and cannot be edited.</p> <ul style="list-style-type: none"> Otherwise, type a name for the port group.
Edge TEP 1 IP (CIDR)	Enter the CIDR for the first NSX Edge TEP. Each node must have a unique Edge TEP 1 IP.
Edge TEP 2 IP (CIDR)	Enter the CIDR for the second NSX Edge TEP. Each node must have a unique Edge TEP 2 IP. The Edge TEP 2 IP must be different than the Edge TEP 1 IP.
Edge TEP Gateway	Enter the IP address for the NSX Edge TEP gateway.
Edge TEP VLAN	Enter the NSX Edge TEP VLAN ID.
First NSX VDS Uplink	<p>Specify an ESXi uplink to map the first NSX Edge node uplink network interface to a physical NIC on the host. The default is <code>uplink1</code>.</p> <p>The information you enter here determines the active uplink on the first VLAN port group used by the NSX Edge node. If you enter <code>uplink3</code>, then <code>uplink3</code> is the active uplink and the uplink you specify for the second NSX VDS uplink is the standby uplink.</p> <p>(cross-cluster only)</p> <p>NOTE For in-cluster NSX Edge cluster expansion, new NSX Edge nodes use the same NSX VDS uplinks as the other Edge nodes hosted on the vSphere cluster.</p>
Second NSX VDS Uplink	<p>Specify an ESXi uplink to map the second NSX Edge node uplink network interface to a physical NIC on the host. The default is <code>uplink2</code>.</p> <p>The information you enter here determines the active uplink on the second VLAN port group used by the NSX Edge node. If you enter <code>uplink4</code>, then <code>uplink4</code> is the active uplink and the uplink you specify for the first NSX VDS uplink is the standby uplink.</p> <p>(cross-cluster only)</p> <p>NOTE For in-cluster NSX Edge cluster expansion, new NSX Edge nodes use the same NSX VDS uplinks as the other Edge nodes hosted on the vSphere cluster.</p>
Add Tier-0 Uplinks	Optional. Click Add Tier-0 Uplinks to add tier-0 uplinks. (Active-Active only)
First Tier-0 Uplink VLAN	Enter the VLAN ID for the first uplink. This is a link from the NSX Edge node to the first uplink network.

Table continued on next page

Continued from previous page

Setting	Description
	(Active-Active only)
First Tier-0 Uplink Interface IP (CIDR)	Enter the CIDR for the first uplink. Each node must have unique uplink interface IPs. (Active-Active only)
Peer IP (CIDR)	Enter the CIDR for the first uplink peer. (EBGP only)
Peer ASN	Enter the ASN for the first uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only)
Second Tier-0 Uplink VLAN	Enter the VLAN ID for the second uplink. This is a link from the NSX Edge node to the second uplink network. (Active-Active only)
Second Tier-0 Uplink Interface IP(CIDR)	Enter the CIDR for the second uplink. Each node must have unique uplink interface IPs. The second uplink interface IP must be different than the first uplink interface IP. (Active-Active only)
Peer IP (CIDR)	Enter the CIDR for the second uplink peer. (EBGP only)
ASN Peer	Enter the ASN for the second uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only)

9. Click **Add More Edge Nodes** to enter configuration settings for additional NSX Edge nodes.

An NSX Edge cluster can contain a maximum of 10 NSX Edge nodes.

- For an NSX Edge cluster with a Tier-0 Service High Availability setting of Active-Active, up to 8 of the NSX Edge nodes can have uplink interfaces.
- For an NSX Edge cluster with a Tier-0 Service High Availability setting of Active-Standby, up to 2 of the NSX Edge nodes can have uplink interfaces.

10. When you are done adding NSX Edge nodes, click **Next**.

11. Review the summary and click **Next**.

SDDC Manager validates the NSX Edge node configuration details.

12. If validation fails, use the **Back** button to edit your settings and try again.

To edit or delete any of the NSX Edge nodes, click the three vertical dots next to an NSX Edge node in the table and select an option from the menu.

13. If validation succeeds, click **Finish** to add the NSX Edge node(s) to the NSX Edge cluster.

You can monitor progress in the Tasks panel.

Remove Edge Nodes from an NSX Edge Cluster

You can remove NSX Edge nodes from an NSX Edge Cluster that you created with SDDC Manager if you need to scale down to meet business needs.

- The NSX Edge cluster must be available in the SDDC Manager inventory and must be **Active**.
- The NSX Edge node must be available in the SDDC Manager inventory.
- The NSX Edge cluster must be hosted on one or more vSphere clusters from the same workload domain.
- The NSX Edge cluster must contain more than two NSX Edge nodes.
- The NSX Edge cluster must not be federated or stretched.
- If the NSX Edge cluster was deployed with a Tier-0 Service High Availability of Active-Active, the NSX Edge cluster must contain two or more NSX Edge nodes with two or more Tier-0 routers (SR component) after the NSX Edge nodes are removed.
- If selected edge cluster was deployed with a Tier-0 Service High Availability of Active-Standby, you cannot remove NSX Edge nodes that are the active or standby node for the Tier-0 router.

For information about deleting an NSX Edge cluster, see [KB 78635](#).

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Click the **Edge Clusters** tab.
4. Click the vertical ellipsis menu for the Edge Cluster you want to expand and select **Shrink Edge Cluster**.
5. Select the Edge node(s) to remove and click **Next**.
6. Review the summary and click **Next**.
SDDC Manager validates the request.
7. If validation fails, use the **Back** button to edit your settings and try again.

NOTE

You cannot remove the active and standby Edge nodes of a Tier-1 router at the same time. You can remove one and then remove the other after the first operation is complete.

8. If validation succeeds, click **Finish** to remove the NSX Edge node(s) from the NSX Edge cluster.
You can monitor progress in the Tasks panel.

Managing Avi Load Balancer in VMware Cloud Foundation

VMware® Avi™ Load Balancer (formerly known as NSX Advanced Load Balancer) allows you to implement centrally-managed distributed load balancing for your application workloads within VMware Cloud Foundation and configure enterprise grade load-balancing, global server load balancing, application security, and container ingress services.

Starting with VMware Cloud Foundation 5.2, you can use SDDC Manager to deploy Avi Load Balancer as a high availability cluster of three VMware® Avi™ Controller instances, each running on a separate VM.

NOTE

Previous version of VMware Cloud Foundation support Avi Load Balancer, but do not deploy or manage the Avi Controller Cluster.

The Avi Controller cluster functions as the control plane and stores and manages all policies related to services and management. All Avi Controllers are deployed in the management domain, even when the Avi Load Balancer is deployed in a VI workload domain.

When you deploy Avi Load Balancer in a workload domain, it is associated with the workload domain's NSX Manager.

NOTE

VMware Cloud Foundation 5.2 does not support deploying Avi Load Balancer on a workload domain that shares its NSX Manager with another workload domain.

VMware Cloud Foundation does not deploy or manage the Service Engine VMs (SEs) that function as the data plane. After deploying the Avi Controller cluster, you can use the Avi Load Balancer UI/API, VMware Aria Automation, or Avi Kubernetes Operator to deploy virtual services for an application, which creates the required Service Engine virtual machines. Service Engines (SEs) are deployed in the workload domain in which the Avi Load Balancer is providing load balancing services. All SEs deployed in a VI workload domain are managed by the Avi Controller that is part of the Avi Load Balancer deployment that is associated with the corresponding NSX instance managing the VI workload domain.

Other important considerations:

- VMware Cloud Foundation does not manage license updates for Avi Load Balancer.
- VMware Cloud Foundation does not manage backing up of Avi Load Balancer configuration database. See the [VMware Avi Load Balancer Documentation](#) for information about configuring scheduled and on-demand backups.
- VMware Cloud Foundation does not manage upgrading Avi Controller Cluster. See the [VMware Avi Load Balancer Documentation](#) for information about upgrading.
- The lifecycle of the Avi Service Engines is managed by each Avi Controller Cluster. You perform updates and upgrades in the Avi Load Balancer web interface, which has checks in place to ensure that you can only upgrade to supported versions.
- If you upgraded from an earlier version of VMware Cloud Foundation and had deployed Avi Load Balancer, SDDC Manager will not be aware of or manage that Avi Load Balancer. You can use SDDC Manager to deploy additional Avi Load Balancers in such an environment.
- In order to use Avi Load Balancer for load balancing services in a vSphere IaaS Control Plane environment, the Avi Load Balancer must be registered with the NSX Manager. See [Registering an Avi Load Balancer cluster with an NSX Manager instance](#).

For more information about how to use and manage Avi Load Balancer see:

- [VMware Avi Load Balancer Documentation](#)
- The [Advanced Load Balancing for VMware Cloud Foundation](#) validated solution

Limitations of Avi Load Balancer in VMware Cloud Foundation

Avi Load Balancer is not supported in workload domains with a shared NSX Manager instance.

Deploy Avi Load Balancer for a Workload Domain

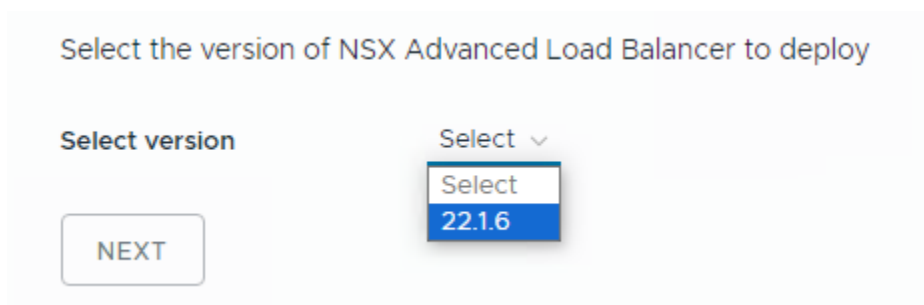
You can deploy Avi Load Balancer for workload domains that do not share their NSX Manager with any other workload domains.

Download the install bundle for a supported version of NSX Advanced Load Balancer. See [Downloading Install Bundles for VMware Cloud Foundation](#).

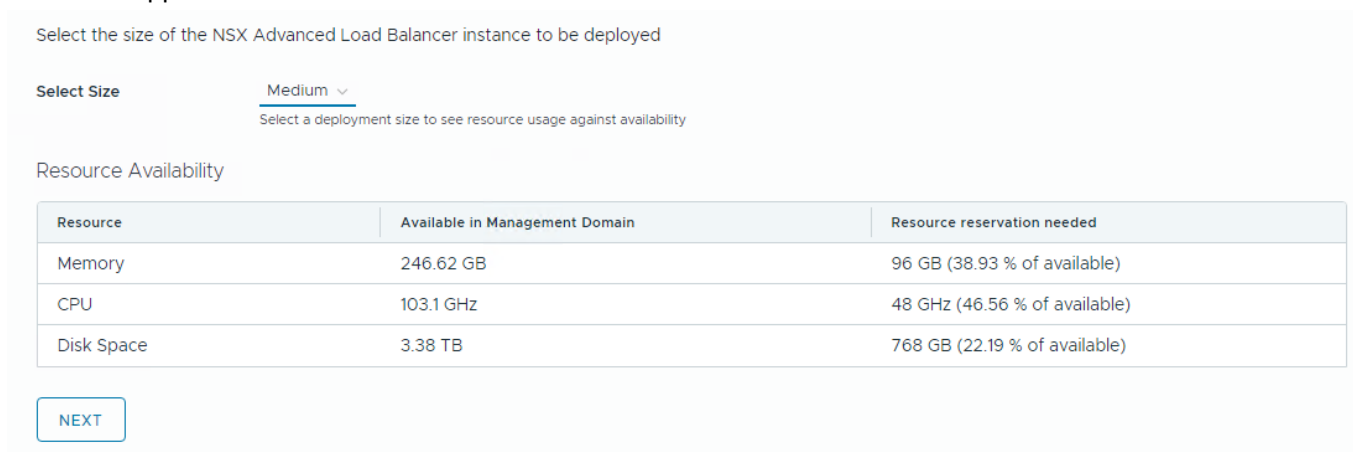
Avi Load Balancer was formerly known as NSX Advanced Load Balancer. The SDDC Manager UI still refers to NSX Advanced Load Balancer.

You cannot deploy Avi Load Balancer on a workload domain that shares its NSX Manager with another workload domain.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Select **Actions** > **Deploy NSX Advanced Load Balancer**.
4. Select the NSX Advanced Load Balancer version and click **Next**.



5. Select the appliance size and click **Next**.



Make sure that the management domain has enough resources for the selected size.

6. Enter the settings for the NSX Advanced Load Balancer Controller cluster and click **Next**.

Administrator password	Enter an administrator password. Although individual VMware Cloud Foundation accounts support different password requirements, it is recommended that you set passwords following a common set of requirements across all accounts: <ul style="list-style-type: none"> • Minimum length: 12 • Maximum length: 20 • At least one lowercase letter, one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ * • Must NOT include: <ul style="list-style-type: none"> – A dictionary word – A palindrome – More than four monotonic character sequences – Three of the same consecutive characters
Node 1 IP Address	Enter the IP address of the first Avi Controller.
Node 2 IP Address	Enter the IP address of the second Avi Controller.
Node 3 IP Address	Enter the IP address of the third Avi Controller.

Table continued on next page

Continued from previous page

Cluster VIP	Enter the Avi Load Balancer Controller cluster IP address. The Avi Load Balancer Controller cluster IP address is a single IP address shared by the Avi Controllers within the cluster. It is the address to which the web interface, CLI commands, and REST API calls are directed. As a best practice, to access the Avi Controller, you must log in to the cluster IP address instead of the IP addresses of individual Avi Controller nodes.
Cluster FQDN	Enter the Avi Controller cluster FQDN. NOTE When creating a service account for the NSX Advanced Load Balancer Controller cluster, VMware Cloud Foundation 5.2 combines the Avi Load Balancer VIP host name and the NSX Manager VIP host name to create the account, <code>svc-<alb hostname>-<nsx hostname></code> . The total characters cannot exceed 32. VCF 5.2.1 automatically truncates the service account name to avoid deployment failures based on account name length.
Cluster Name	Enter a name for the Avi Controller cluster.

7. Click **Start Deployment**.

You can monitor the deployment in the Tasks panel.

Task	Subtask	Task Status	Last Occurrence
Adding NSX Advanced Load Balancer Cluster nsx-alb	Deploy NSX Advanced Load Balancer on NSX. This includes OVA upload, VM creation and wait for cluster HA.	<div style="width: 42%;"><div style="background-color: #0070c0; height: 10px;"></div></div> 42%	5/11/24, 1:05 AM
Download BUNDLE - NSX_ALB:22.16-23390967		✔ Successful	5/10/24, 11:58 PM

After the Avi Load Balancer Controller cluster deploys successfully, you can access the web interface from the **Services** tab for the workload domain by clicking the NSX Advanced Load Balancer link.

MANAGEMENT ✔ ACTIVE Version : 5.2.0.0

Summary Services Updates Update History

VMware Cloud Foundation Components

Component
vCenter Server vcenter-1.vrack.vsphere.local ↗
NSX Cluster vip-nsx-mgmt.vrack.vsphere.local ↗
NSX Advanced Load Balancer vipa.vrack.vsphere.local ↗

You can manage the Avi Load Balancer Controller cluster administrator password and certificate using the SDDC Manager UI.

- [Managing Passwords in VMware Cloud Foundation](#)
- [Managing Certificates in VMware Cloud Foundation](#)

Remove Avi Load Balancer from a Workload Domain

If you deployed an Avi Load Balancer to a workload domain and you no longer need it, you can remove it from the workload domain. If a workload domain includes an Avi Load Balancer, you cannot delete the workload domain until you remove the Avi Load Balancer.

The Avi Load Balancer must not be in use.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Select **Actions** › **Remove NSX Advanced Load Balancer**.

Add Cluster
Add Edge Cluster
Add AVNs
Update Licenses
Rename Domain
Remove NSX Advanced Load Balancer

4. Click **Remove** to confirm.

Remove NSX Advanced Load Balancer



Are you sure you want to remove NSX Advanced Load Balancer from this workload domain?

CANCEL

REMOVE

Deploying Application Virtual Networks in VMware Cloud Foundation

Before you can deploy VMware Aria Suite components or implement the Identity and Access Management for VMware Cloud Foundation validated solution, you must deploy Application Virtual Networks in the management domain.

An Application Virtual Network (AVN) is a software-defined networking concept based on NSX that allows the hosting of management applications on NSX segments. In NSX, segments are virtual layer-2 domains.

You can create overlay-backed NSX segments or VLAN-backed NSX segments. Both options create two NSX segments (Region-A and X-Region) on the NSX Edge cluster deployed in the default management vSphere cluster. Those NSX segments are used when you deploy the VMware Aria Suite products. Region-A segments are local instance NSX segments and X-Region segments are cross-instance NSX segments.

IMPORTANT

You cannot create AVNs if NSX for the management domain is part of an NSX Federation. See [Working with NSX Federation in VMware Cloud Foundation](#).

Overlay-Backed NSX Segments

Overlay-backed segments provide flexibility for workload placement by removing the dependence on traditional data center networks. Using overlay-backed segments improves the security and mobility of management applications and reduces the integration effort with existing networks. Overlay-backed segments are created in an overlay transport zone.

In an overlay-backed segment, traffic between two VMs on different hosts but attached to the same overlay segment have their layer-2 traffic carried by a tunnel between the hosts. NSX instantiates and maintains this IP tunnel without the need for any segment-specific configuration in the physical infrastructure. As a result, the virtual network infrastructure is decoupled from the physical network infrastructure. That is, you can create segments dynamically without any configuration of the physical network infrastructure.

VLAN-Backed NSX Segments

VLAN-backed segments leverage the physical data center networks to isolate management applications, while still taking advantage of NSX to manage these networks. VLAN-backed network segments ensure the security of management applications without requiring support for overlay networking. VLAN-backed segments are created in a VLAN transport zone.

A VLAN-backed segment is a layer-2 broadcast domain that is implemented as a traditional VLAN in the physical infrastructure. This means that traffic between two VMs on two different hosts but attached to the same VLAN-backed segment is carried over a VLAN between the two hosts. The resulting constraint is that you must provision an appropriate VLAN in the physical infrastructure for those two VMs to communicate at layer-2 over a VLAN-backed segment.

VMware Aria Suite Components and NSX Segments

When you deploy the VMware Aria Suite components, they use the NSX segments that you created.

VMware Aria Suite Component	NSX Segment
VMware Aria Operations for Logs	Region-A
VMware Aria Operations Manager	X-Region
Workspace ONE Access	X-Region
VMware Aria Automation	X-Region
VMware Aria Suite Lifecycle	X-Region

Identity and Access Management for VMware Cloud Foundation

See [Identity and Access Management for VMware Cloud Foundation](#) for more information about how that validated solution uses Application Virtual Networks.

Deploy Overlay-Backed NSX Segments

Create overlay-backed NSX segments, also known as Application Virtual Networks (AVNs), for use with VMware Aria Suite components.

Create an NSX Edge cluster for Application Virtual Networks, using the recommended settings, in the default management vSphere cluster. See [Deploy an NSX Edge Cluster](#).

This procedure describes creating overlay-backed NSX segments. If you want to create VLAN-backed NSX segments instead, see [Deploy VLAN-Backed NSX Segments](#).

1. In the navigation page, click **Inventory** › **Workload Domains**.
2. Click on the management domain.
3. Select **Actions** › **Add AVNs**.
4. Select **Overlay-backed network segment** and click **Next**.
5. Select an NSX Edge cluster and a Tier-1 gateway.
6. Enter information for each of the NSX segments (Region-A and X-Region):

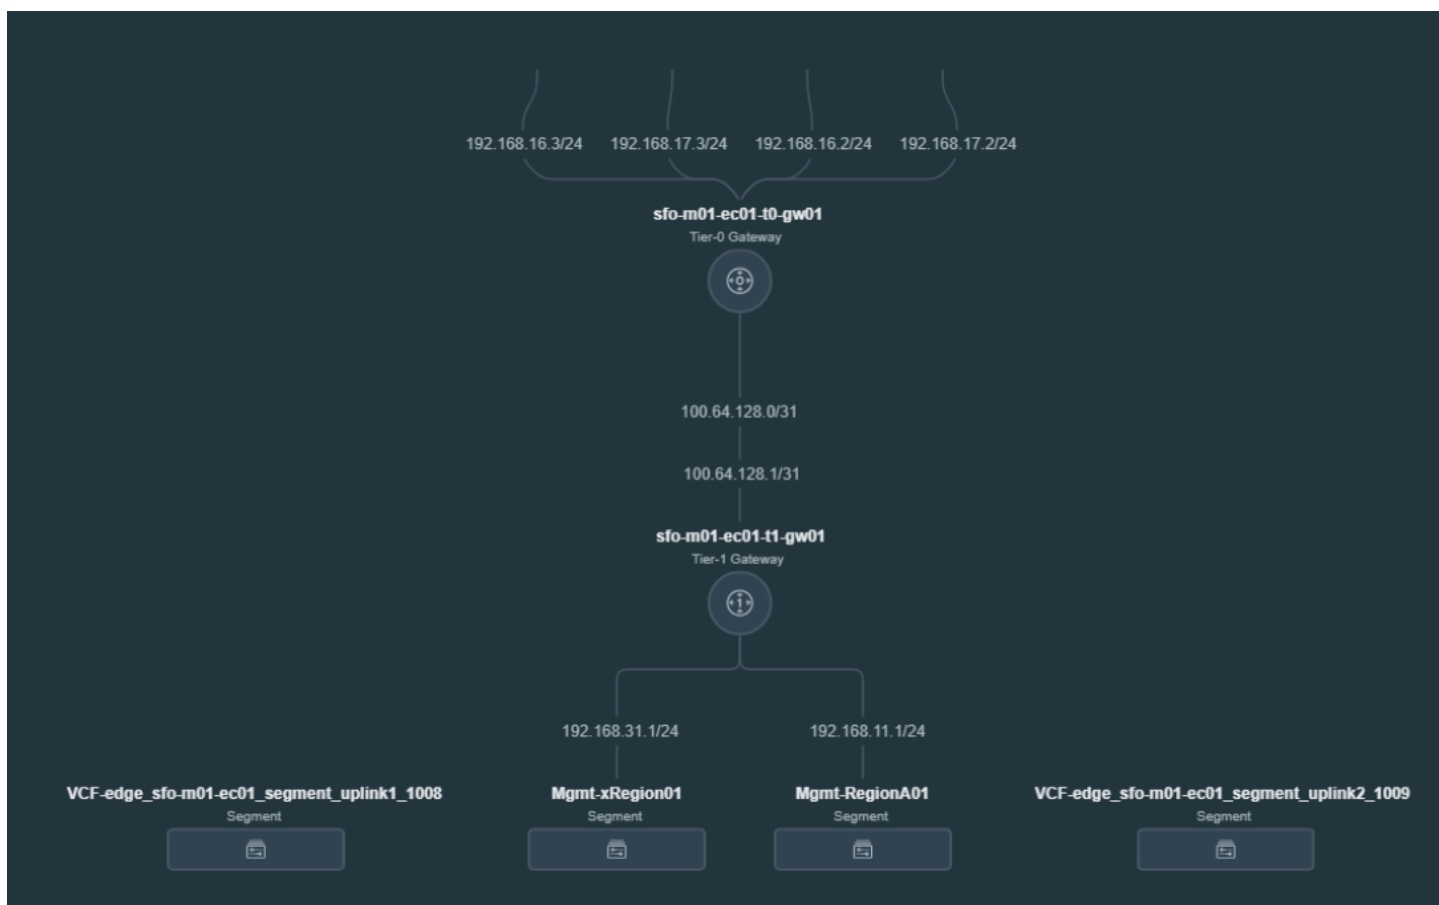
Option	Description
Name	Enter a name for the NSX segment. For example, <code>Mgmt-RegionA01</code> .
Subnet	Enter a subnet for the NSX segment.
Subnet mask	Enter a subnet mask for the NSX segment.
Gateway	Enter a gateway for the NSX segment.
MTU	Enter an MTU for the NSX segment.

7. Click **Validate Settings** and then click **Next**.

If validation does not succeed, verify and update the information you entered for the NSX segments and click **Validate Settings** again.

8. Review the settings and click **Finish**.

Example Network Topology for Overlay-Backed NSX Segments



Deploy VLAN-Backed NSX Segments

Create VLAN-backed NSX segments, also known as Application Virtual Networks (AVNs), for use with VMware Aria Suite components.

Create an NSX Edge cluster for Application Virtual Networks, using the recommended settings, in the default management vSphere cluster. See [Deploy an NSX Edge Cluster](#).

You must have an available VLAN ID for each NSX segment.

This procedure describes creating VLAN-backed NSX segments. If you want to create overlay-backed NSX segments instead, see [Deploy Overlay-Backed NSX Segments](#).

1. In the navigation page, click **Inventory** > **Workload Domains**.
2. Click on the management domain.
3. Select **Actions** > **Add AVNs**.
4. Select **VLAN-backed network segment** and click **Next**.
5. Select an NSX Edge cluster.
6. Enter information for each of the NSX segments (Region-A and X-Region):

Option	Description
Name	Enter a name for the NSX segment. For example, Mgmt-RegionA01.
Subnet	Enter a subnet for the NSX segment.

Table continued on next page

Continued from previous page

Option	Description
Gateway	Enter a gateway for the NSX segment.
MTU	Enter an MTU for the NSX segment.
VLAN ID	Enter the VLAN ID for the NSX segment.

7. Click **Validate Settings** and then click **Next**.

If validation does not succeed, verify and update the information you entered for the NSX segments and click **Validate Settings** again.

8. Review the settings and click **Finish**.

Example Network Topology for VLAN-Backed NSX Segments



VMware Cloud Foundation with VMware Tanzu

VMware Cloud Foundation™ with VMware Tanzu™ enables you to deploy and operate the compute, networking, and storage infrastructure for vSphere IaaS Control Plane workloads. vSphere IaaS Control Plane transforms vSphere to a platform for running Kubernetes workloads natively on the hypervisor layer.

When enabled on a vSphere cluster, vSphere IaaS Control Plane provides the capability to run Kubernetes workloads directly on ESXi hosts and to create upstream Kubernetes clusters within dedicated resource pools. vSphere IaaS Control Plane can also be enabled on the management domain default cluster.

NOTE

Starting with vSphere 8.0 Update 3, vSphere with Tanzu was renamed to vSphere IaaS Control Plane.

You validate the underlying infrastructure for vSphere IaaS Control Plane from the SDDC Manager UI and then complete the deployment in the vSphere Client. The SDDC Manager UI refers to the vSphere IaaS Control Plane functionality as Kubernetes - Workload Management.

The [Developer Ready Infrastructure for VMware Cloud Foundation](#) validated solution provides design, implementation, and operational guidance for a workload domain that runs vSphere with Tanzu workloads in the Software-Defined Data Center (SDDC).

For more information about vSphere IaaS Control Plane, see [What Is vSphere IaaS Control Plane?](#)

Enable Workload Management

With Workload Management, you validate the underlying infrastructure for vSphere IaaS Control Plane. You then complete the deployment using the vSphere Client.

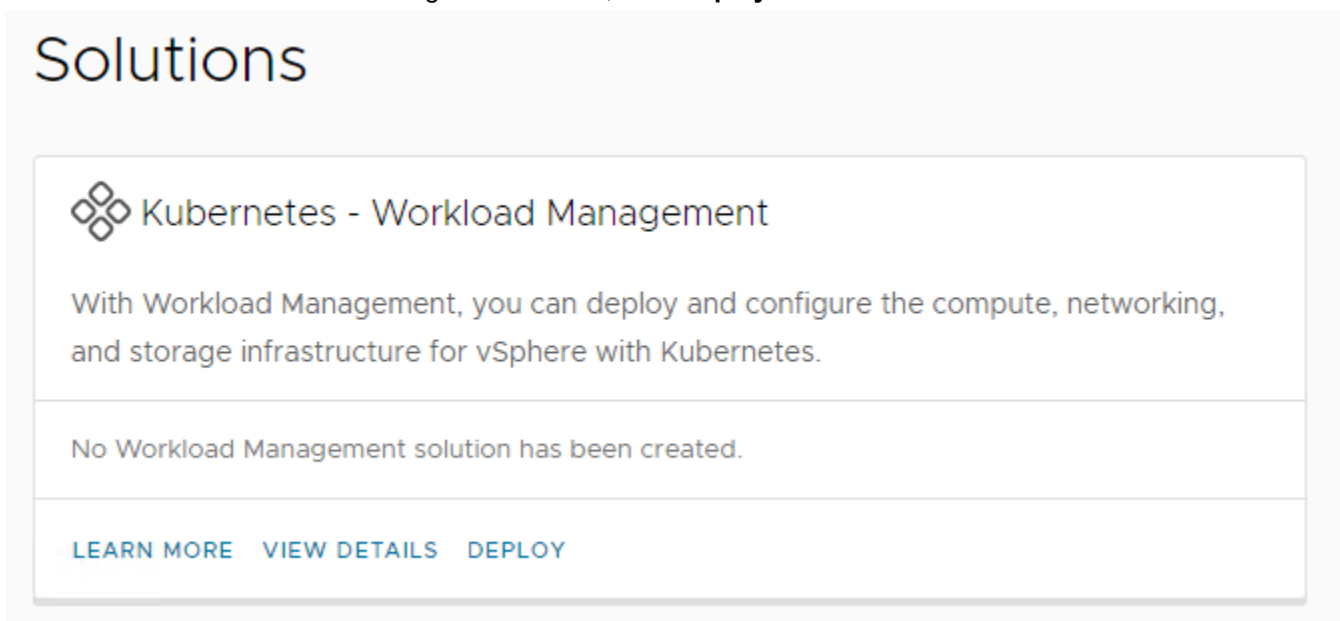
- A VI workload domain must be deployed.

NOTE


If you deployed VMware Cloud Foundation with a consolidated architecture, you can enable Workload Management on the management domain.

- An Workload Management ready NSX Edge cluster must be deployed on the workload domain. You must select Workload Management on the Use Case page of the Add Edge Cluster wizard. See step 6 in [Deploy an NSX Edge Cluster](#).
- All hosts in the vSphere cluster for which you enable Workload Management must be licensed for vSphere IaaS Control Plane.
- Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.
- The following IP address subnets must be defined:
 - A non-routable subnet for pod networking, minimum of a /22 subnet.
 - A non-routable subnet for Service IP addresses, minimum of a /24 subnet
 - A routable subnet for ingress, minimum of a /27 subnet
 - A routable subnet for egress, minimum of a /27 subnet
- In order to use Avi Load Balancer for load balancing services in a vSphere IaaS Control Plane environment, the Avi Load Balancer must be registered with the NSX Manager. See [Registering an Avi Load Balancer cluster with an NSX Manager instance](#).

1. In the navigation pane, click **Solutions**.
2. In the Kubernetes - Workload Management section, click **Deploy**.



Solutions

 **Kubernetes - Workload Management**

With Workload Management, you can deploy and configure the compute, networking, and storage infrastructure for vSphere with Kubernetes.

No Workload Management solution has been created.

[LEARN MORE](#) [VIEW DETAILS](#) [DEPLOY](#)

3. Review the Workload Management prerequisites, click **Select All**, and click **Begin**.
4. Select the workload domain associated with the vSphere cluster where you want to enable Workload Management. The Workload Domain drop-down menu displays all Workload Management ready workload domains, including the management domain.

vSphere clusters in the selected workload domain that are compatible with Workload Management are displayed in the Compatible section. Incompatible clusters are displayed in the Incompatible section, along with the reason for the incompatibility. If you want to get an incompatible cluster to a usable state, you can exit the Workload Management deployment wizard while you resolve the issue.

- From the list of compatible clusters on the workload domain, select the cluster where you want to enable Workload Management and click **Next**.
- On the Validation page, wait for validation to complete successfully and click **Next**.

The following validations are performed.

- vCenter Server validation (vCenter Server credentials, vSphere cluster object, and version)
- Network validation (NSX Manager credentials and version)
- Compatibility validation (vSphere cluster and content library)

- On the Review page, review your selections and click **Complete in vSphere**. You are automatically redirected to the vSphere Client.

Follow the deployment wizard within the vSphere Client to complete the Workload Management deployment and configuration steps.

View Workload Management Cluster Details

The Workload Management page displays clusters with Workload Management. The status of each cluster, number of hosts in the cluster, and associated workload domain is also displayed.

- In the navigation pane, click **Solutions**.
- In the Kubernetes - Workload Management section, click **View Details**.
- Click vSphere Workload Management Clusters to see cluster details in vSphere.

Update Workload Management License

Once you enable Workload Management on a cluster, you must assign a Tanzu edition license to the cluster before the evaluation license expires.

You must have added a VMware Tanzu license key to the Cloud Foundation license inventory. See [Add a Component License Key in the SDDC Manager UI](#).

- In the navigation pane, click **Solutions**.
- Click the dots to the left of the cluster for which you want to update the license and click **Update Workload Management license**.
- Select the appropriate license and click **Apply**.
After the license update processing is completed, the Workload Management page is displayed. The task panel displays the licensing task and its status.

VMware Aria Suite Lifecycle in VMware Cloud Foundation mode

When you deploy VMware Aria Suite Lifecycle from the SDDC Manager UI, VMware Cloud Foundation mode is enabled in VMware Aria Suite Lifecycle, and the behavior of VMware Aria Suite Lifecycle is aligned with the VMware Cloud Foundation architecture.

VMware Aria Suite Lifecycle in VMware Cloud Foundation mode introduces the following features:

- Automatic load balancer configuration. Load balancer preparation and configuration are no longer a prerequisite when you use VMware Aria Suite Lifecycle to deploy or perform a cluster expansion on Workspace ONE Access,

VMware Aria Operations, or VMware Aria Automation. Load balancer preparation and configuration take place as part of the deploy or expand operation.

- Automatic infrastructure selection in the VMware Aria Suite Lifecycle deployment wizards. When you deploy a VMware Aria Suite product through VMware Aria Suite Lifecycle, infrastructure objects such as clusters and networks are pre-populated. They are fixed and cannot be changed to ensure alignment with the VMware Cloud Foundation architecture.
- Cluster deployment for a new environment. You can deploy VMware Aria Operations for Logs, VMware Aria Operations, or VMware Aria Automation in clusters. You can deploy Workspace ONE Access either as a cluster or a single node. If you deploy Workspace ONE Access as a single node, you can expand it to a cluster later.
- Consistent Bill Of Materials (BOM). VMware Aria Suite Lifecycle in VMware Cloud Foundation mode only displays product versions that are compatible with VMware Cloud Foundation to ensure product interoperability.
- Inventory synchronization between VMware Aria Suite Lifecycle and SDDC Manager. VMware Aria Suite Lifecycle can detect changes made to VMware Aria Suite products and update its inventory through inventory synchronization. When VMware Cloud Foundation mode is enabled in VMware Aria Suite Lifecycle, inventory synchronization in VMware Aria Suite Lifecycle also updates SDDC Manager's inventory to get in sync with the current state of the system.
- Product versions. VMware Cloud Foundation supports flexible VMware Aria Suite upgrades. You can upgrade VMware Aria Suite products as new versions become available in VMware Aria Suite Lifecycle. VMware Aria Suite Lifecycle will only allow upgrades to compatible and supported versions of VMware Aria Suite products.
- Resource pool and advanced properties. The resources in the Resource Pools under the Infrastructure Details are blocked by the VMware Aria Suite Lifecycle UI, so that the VMware Cloud Foundation topology does not change. Similarly, the Advanced Properties are also blocked for all products except for Remote Collectors. VMware Aria Suite Lifecycle also auto-populates infrastructure and network properties by calling VMware Cloud Foundation deployment API.
- Federal Information Processing Standard (FIPS) compliance.
- Watermark.

VMware Aria Suite Lifecycle Implementation

You deploy VMware Aria Suite Lifecycle in VMware Cloud Foundation mode by using SDDC Manager. After that, you perform the necessary post-deployment configurations.

- Download the VMware Software Install Bundle for VMware Aria Suite Lifecycle from the VMware Depot to the local bundle repository. See [Downloading Install Bundles for VMware Cloud Foundation](#).
- Allocate an IP address for the VMware Aria Suite Lifecycle virtual appliance on the cross-instance NSX segment and prepare both forward (A) and reverse (PTR) DNS records.
- Allocate an IP address for the NSX standalone Tier-1 Gateway on the cross-instance NSX segment. This address is used for the service interface of the standalone NSX Tier 1 Gateway created during the deployment. The Tier 1 Gateway is used for load-balancing of specific VMware Aria Suite products and Workspace ONE Access.
- Ensure you have enough storage capacity:
 - Required storage: 178 GB
 - Virtual disk provisioning: Thin
- Verify that the management domain vCenter Server is operational.
- Verify that NSX Manager is operational.
- Verify the **Prerequisite Checklist** sheet in the *Planning and Preparation Workbook*.

By default, VMware Cloud Foundation uses NSX to create NSX segments and deploys VMware Aria Suite Lifecycle and the VMware Aria Suite products to these NSX segments. Starting with VMware Cloud Foundation 4.3, NSX segments are no longer configured during the management domain bring-up process, but instead are configured using the SDDC Manager UI. The new process offers the choice of using either overlay-backed or VLAN-backed segments. See [Deploying Application Virtual Networks in VMware Cloud Foundation](#).

VMware Aria Suite Lifecycle runs in VMware Cloud Foundation mode, the integration ensures awareness between the two components. You launch the deployment of VMware Aria Suite products from the SDDC Manager UI and are redirected to the VMware Aria Suite Lifecycle UI where you complete the deployment process.

Deploy VMware Aria Suite Lifecycle

You deploy the VMware Aria Suite Lifecycle in VMware Cloud Foundation mode by using the SDDC Manager UI.

1. In the navigation pane, click **Administration** > **VMware Aria Suite**.
2. Click **Deploy**.
3. Review and verify the prerequisites.
Click each prerequisite check box and then click **Begin**.
4. On the **Network Settings** page, review the settings and click **Next**.
5. On the **Virtual Appliance Settings** page, enter the settings and click **Next**.

Setting	Description
Virtual Appliance: FQDN	The FQDN for the VMware Aria Suite Lifecycle virtual appliance. NOTE The reverse (PTR) DNS record of this fully qualified domain name is used as the IP address for the virtual appliance.
NSX Tier 1 Gateway: IP Address	A free IP Address within the cross-instance virtual network segment. NOTE Used to create a service interface on the NSX Tier 1 Gateway, where VMware Cloud Foundation automatically configures the load-balancer for the VMware Aria Suite.
System Administrator	Create and confirm the password for the VMware Aria Suite Lifecycle administrator account, vcfadmin@local . The password created is the credential that allows SDDC Manager to connect to VMware Aria Suite Lifecycle. NOTE When VMware Aria Suite Lifecycle is deployed by SDDC Manager it is enabled for VMware Cloud Foundation mode. As a result, the administrator account for is vcfadmin@local instead of admin@local .
SSH Root Account	Create and confirm a password for the VMware Aria Suite Lifecycle virtual appliance root account.

6. On the **Review Summary** page, review the installation configuration settings and click **Finish**.
SDDC Manager validates the values and starts the deployment.

The VMware Aria Suite page displays the following message: `Deployment in progress`.

If the deployment fails, this page displays a deployment status of `Deployment failed`. In this case, you can click **Restart Task** or **Rollback**.

7. **(Optional)** To view details about the individual deployment tasks, in the **Tasks** panel at the bottom, click each task.

Replace the Certificate of the VMware Aria Suite Lifecycle Instance

To establish a trusted connection to VMware Aria Suite Lifecycle, you replace the SSL certificate on the appliance by using the SDDC Manager UI.

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the **Workload Domain** page, from the table, in the domain column click the management domain.
3. On the domain summary page, click the **Certificates** tab.
4. From the table, select the check box for the VMware Aria Suite Lifecycle resource type, and click **Generate CSRs**.
5. On the **Details** page, enter the following settings and click **Next**.

Settings	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

6. On the **Subject Alternative Name** page, leave the default SAN and click **Next**.
7. On the **Summary** page, click **Generate CSRs**.
8. After the successful return of the operation, click **Generate signed certificates**.
9. In the **Generate Certificates** dialog box, from the **Select Certificate Authority** drop-down menu, select **Microsoft**.
10. Click **Generate certificates**.
11. After the successful return of the operation, click **Install certificates**.
Wait for the successful return of the operation.

Configure Data Center and vCenter Server in VMware Aria Suite Lifecycle

Before you can create a global environment for product deployments, you must add a cross-instance data center and the associated management domain vCenter Server to VMware Aria Suite Lifecycle.

You add the cross-instance data center, and the associated management domain vCenter Server for the deployment of the global components, such as the clustered Workspace ONE Access.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).

2. On the **My Services** page, click **Lifecycle Operations**.
3. In the navigation pane, click **Datacenters**.
4. Click **Add datacenter**, enter the values for the global data center, and click **Save**.

Setting	Value
Datacenter name	Name for cross-instance datacenter
Use custom location	Deactivated
Location	Location of datacenter

5. Add the management domain vCenter Server to the global data center.
 - a) On the **Datacenters** page, expand the global data center and click **Add vCenter**.
 - b) Enter the management domain vCenter Server information and click **Validate**.

Setting	Value
vCenter name	Enter a name for the vCenter Server
vCenter FQDN	Enter the FQDN of the vCenter Server
vCenter credentials	Select the <code><management_vcenter_name>-<uid></code> credential. For example: vcenter-1-35214fac-caeb-4062-a184-350344e30c7f.
vCenter type	Management

6. After the successful vCenter Server validation, click **Save**.
7. In the navigation pane, click **Requests** and verify that the state of the **vCenter data collection request** is Completed.

Workspace ONE Access Implementation

Workspace ONE Access provides identity and access management services for the VMware Aria Suite of products. You use VMware Aria Suite Lifecycle to deploy a Workspace ONE Access instance. You then perform the necessary post-deployment configurations and customization. VMware Cloud Foundation supports both standard and clustered Workspace ONE Access deployments.

- Download the installation binary directly from VMware Aria Suite Lifecycle. See "Configure Product Binaries" in the *VMware Aria Suite Lifecycle Installation, Upgrade, and Management Guide* for the version of [VMware Aria Suite Lifecycle](#) listed in the VMware Cloud Foundation BOM.
- Allocate IP addresses:

Standard Deployment	Clustered Deployment
One IP address from the cross-instance NSX segment and prepare both forward (A) and reverse (PTR) DNS records.	Five IP addresses from the cross-instance NSX segment and prepare both forward (A) and reverse (PTR) DNS records. <ul style="list-style-type: none"> • Three IP addresses for the clustered Workspace ONE Access instance. • One IP address for the embedded Postgres database for the Workspace ONE Access instance.

Table continued on next page

Continued from previous page

Standard Deployment	Clustered Deployment
	<ul style="list-style-type: none"> One IP address for the NSX external load balancer virtual server for clustered Workspace ONE Access instance.

- Ensure you have enough storage capacity:
 - Required storage per node: 100 GB
 - Virtual disk provisioning: Thin
- Verify that the management domain vCenter Server is operational.
- Verify that the cross-instance NSX segment is available.
- Verify that the NSX Manager is operational.
- Verify the **Prerequisite Checklist** sheet in the *Planning and Preparation Workbook*.
- Verify that required Active Directory bind service account is created.

Verify that required Active Directory security groups are created.

- Download the `CertGenVVS` tool and generate the signed certificate for the Workspace ONE Access instance. See [KB 85527](#).

Import the Workspace ONE Access Certificate to VMware Aria Suite Lifecycle

To prepare VMware Aria Suite Lifecycle for deploying Workspace ONE Access, you must generate an SSL certificate using the PowerShell module for VMware Validated Solutions and add the certificate to the VMware Aria Suite Lifecycle locker.

- Verify that a Microsoft Certificate Authority is available for the environment.
- Install the [PowerShell module for VMware Validated Solutions](#) together with the supporting modules to request an SSL certificate from your Microsoft Certificate Authority.
- Verify that you have OpenSSL 3.0 or later installed on the system that will run the PowerShell module. The [OpenSSL Wiki](#) has a list of third-party pre-compiled binaries for Microsoft Windows.

This procedure uses the PowerShell Module for VMware Validated Solutions to generate the required certificates from a Microsoft Active Directory Certificate Services. However, the module also supports generating certificate signing requests (CSRs) for third party certificate authorities for import to the VMware Aria Suite Lifecycle locker.

- Generate an SSL certificate using the PowerShell module for VMware Validated Solutions.
 - Start PowerShell.
 - Replace the sample values in the variables below and run the commands in the PowerShell console.

```
$commonName = "xint-idm01.rainpole.io"

$subjectAltNames = "xint-idm01.rainpole.io, xint-idm01a.rainpole.io, xint-idm01b.rainpole.io, xint-cidm01c.rainpole.io"

$encryptionKeySize = 2048

$certificateExpiryDays = 730

$orgName = "rainpole"

$orgUnitName = "Platform Engineering"

$orgLocalityName = "San Francisco"

$orgStateName = "California"
```

```

$orgCountryCode = "US"

$caType = "msca"
$caFqdn = "rpl-ad01.rainpole.io"
$caUsername = "Administrator"
$caPassword = "VMw@re1!"
$caTemplate = "VMware"

$outputPath = ".\certificates\"
$csrFilePath = Join-Path $outputPath "$commonName.csr"
$keyFilePath = Join-Path $outputPath "$commonName.key"
$certFilePath = Join-Path $outputPath "$commonName.crt"
$rootCaFilePath = Join-Path $outputPath "$caFqdn-rootCa.pem"

```

c) Perform the configuration by running the command in the PowerShell console.

```

Invoke-GeneratePrivateKeyAndCsr -outDirPath $outputPath -commonName $commonName
-subjectAlternativeNames $subjectAltNames -keySize $encryptionKeySize
-expireDays $certificateExpiryDays -organization $orgName -organizationUnit
$orgUnitName -locality $orgLocalityName -state $orgStateName -country
$orgCountryCode

Invoke-RequestSignedCertificate -caFqdn $caFqdn -csrFilePath $csrFilePath
-outDirPath $outputPath -certificateAuthority $caType -username $caUsername
-password $caPassword -certificateTemplate $caTemplate -getCArootCert

Invoke-GenerateChainPem -outDirPath $outputPath -keyFilePath $keyFilePath
-crtFilePath $certFilePath -rootCaFilePath $rootCaFilePath

```

2. Add the generated SSL certificate to the VMware Aria Suite Lifecycle locker.

- a) Log in to VMware Aria Suite Lifecycle at https://<aria_suite_lifecycle_fqdn> as `vcfadmin@local`.
- b) On the **My services** page, click **Locker**.
- c) In the navigation pane, click **Certificates**.
- d) On the **Certificates** page, click **Import**.
- e) On the **Import certificate** page, enter a name for the Workspace ONE Access certificate according to your *VMware Cloud Foundation Planning and Preparation Workbook*.
- f) Click **Browse file**, navigate to the Workspace ONE Access certificate file (`.pem`), and click **Open**.
- g) On the **Import certificate** page, click **Import**.

Add Workspace ONE Access Passwords to VMware Aria Suite Lifecycle

To enable life cycle management and configuration management, you set the passwords for the VMware Aria Suite Lifecycle cross-instance environment administrator account and for the Workspace ONE Access administrator and configuration administrator accounts.

You add the following passwords for the corresponding local administrative accounts.

Setting	Value for Global Environment Administrator	Value for Local Administrator	Value for Local Configuration Administrator	Value for Appliance Root User
Password alias	global-env-admin	xint-wsa-admin	xint-wsa-configadmin	xint-wsa-root
Password	<i>global_env_admin_password</i>	<i>xint_wsa_admin_password</i>	<i>xint_wsa_configadmin_password</i>	<i>xint_wsa_root_password</i>
Confirm password	<i>global_env_admin_password</i>	<i>xint-wsa_admin_password</i>	<i>xint_wsa_configadmin_password</i>	<i>xint_wsa_root_password</i>
Password description	VMware Aria Suite Lifecycle global environment default password Used for Workspace ONE Access appliance sshuser.	Workspace ONE Access administrator	Workspace ONE Access configuration administrator	Workspace ONE Access root user

NOTE

You do not need to provide a user name when adding passwords. You can leave the **User Name** field blank when configuring settings.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Locker**.
3. In the navigation pane, click **Passwords**.
4. On the **Passwords** page, click **Add**.
5. On the **Add password** page, configure the settings and click **Add**.
6. Repeat this procedure for all the remaining credentials.

Deploy a Standard Workspace ONE Access Instance Using VMware Aria Suite Lifecycle

To provide identity and access management services to the cross-instance SDDC components, you create a global environment in VMware Aria Suite Lifecycle in which you deploy a standard Workspace ONE Access instance.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Lifecycle Operations**.
3. On the **Dashboard** page, click **Create environment**.
4. On the **Create environment** page, configure the settings and click **Next**.

Setting	Value
Install Identity Manager	Selected
Default password	global-env-admin
Datacenter	Select the cross-instance datacenter.
JSON configuration	Deactivated

Table continued on next page

Continued from previous page

Setting	Value
Join the VMware customer experience improvement program	Selected

5. On the **Select product** page, select the check box for **VMware Identity Manager**, configure these values, and click **Next**.

Setting	Value
Installation type	New install
Version	Select a version. VMware Aria Suite Lifecycle will only display supported versions.
Deployment type	Standard

6. On the **Accept license agreements** page, scroll to the bottom and accept the license agreement, and then click **Next**.
7. On the **Certificate** page, from the **Select certificate** drop-down menu, select the *Workspace One Access* certificate, and click **Next**.
8. On the **Infrastructure** page, verify and accept the default settings, and click **Next**.
9. On the **Network** page, verify and accept the default settings, and click **Next**.
10. On the **Products** page, configure the deployment properties of Workspace ONE Access and click **Next**.
- a) In the **Product properties** section, configure the settings.

Setting	Value
Certificate	<i>Workspace One Access</i>
Node size	Medium (VMware Aria Automation recommended size)
Admin password	Select the <i>xint-wsa-admin</i>
Default configuration admin email	Enter a default email.
Default configuration admin user name	<i>configadmin</i>
Default configuration admin password	Select the <i>xint-wsa-configadmin</i>
Sync group members	Selected

- b) In the **Components** section, configure the primary node.

Setting	Value for vidm-primary
VM Name	Enter a VM Name for vidm-primary.
FQDN	Enter the FQDN for vidm-primary
IP address	Enter the IP Address for vidm-primary.

- c) Click advanced configuration and click **Select Root Password**.
- d) Select *xint-wsa-root* and click **Save**.
11. On the **Precheck** page, click **Run precheck**.
12. On the **Manual validations** page, select the **I took care of the manual steps above and am ready to proceed** check box and click **Run precheck**.

13. Review the validation report, remediate any errors, and click **Re-run precheck**.
14. Wait for all prechecks to complete with `Passed` messages and click **Next**.
15. On the **Summary** page, review the configuration details. To back up the deployment configuration, click **Export configuration**.
16. To start the deployment, click **Submit**.
The **Request details** page displays the progress of deployment.
17. Monitor the steps of the deployment graph until all stages become `Completed`.

Deploy Clustered Workspace ONE Access Instance Using VMware Aria Suite Lifecycle

To provide identity and access management services to the cross-instance SDDC components, you create a global environment in VMware Aria Suite Lifecycle in which you deploy a 3-node clustered Workspace ONE Access instance.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Lifecycle Operations**.
3. On the **Dashboard** page, click **Create environment**.
4. On the **Create environment** page, configure the settings and click **Next**.

Setting	Value
Install Identity Manager	Selected
Default password	global-env-admin
Datacenter	Select the cross-instance datacenter.
JSON configuration	Deactivated
Join the VMware customer experience improvement program	Selected

5. On the **Select product** page, select the check box for **VMware Identity Manager**, configure these values, and click **Next**.

Setting	Value
Installation type	New install
Version	Select a version. VMware Aria Suite Lifecycle will only display supported versions.
Deployment type	Cluster

6. On the **Accept license agreements** page, scroll to the bottom and accept the license agreement, and then click **Next**.
7. On the **Certificate** page, from the **Select certificate** drop-down menu, select the *Clustered Workspace One Certificate*, and click **Next**.
8. On the **Infrastructure** page, verify and accept the default settings, and click **Next**.
9. On the **Network** page, verify and accept the default settings, and click **Next**.
10. On the **Products** page, configure the deployment properties of clustered Workspace ONE Access and click **Next**.
 - a) In the **Product properties** section, configure the settings.

Setting	Value
Certificate	<i>Workspace One Access</i>
Node size	Medium (VMware Aria Automation recommended size)
Admin password	Select the <i>xint-wsa-admin</i>
Default configuration admin email	Enter a default email.
Default configuration admin user name	<i>configadmin</i>
Default configuration admin password	Select the <i>xint-wsa-configadmin</i>
Sync group members	Selected

b) In the **Cluster Virtual IP** section, click **Add Load Balancer** and configure its settings.

Setting	Value
Controller Type	VMware Cloud Foundation managed NSX-T
Load Balancer IP	Use the IP address from your <i>VMware Cloud Foundation Planning and Preparation Workbook</i> .
Load Balancer FQDN	Use the FQDN from your <i>VMware Cloud Foundation Planning and Preparation Workbook</i> .

c) In the **Cluster VIP FQDN** section, configure the settings.

Setting	Value
Controller Type	Select VMware Cloud Foundation managed NSX-T from the drop-down menu.
FQDN	Select the Load Balancer FQDN from the drop-down menu.
Locker certificate	Clustered Workspace ONE Access Certificate
Database IP address	Enter the IP address for the embedded Postgres database. NOTE The IP address must be a valid IP address for the cross-instance NSX segment.

d) In the **Components** section, configure the three cluster node.

Setting	Value for vidm-primary	Value for vidm-secondary-1	Value for vidm-secondary-2
VM Name	Enter a VM Name for vidm-primary.	Enter a VM Name for vidm-secondary-1.	Enter a VM Name for vidm-secondary-2.

Table continued on next page

Continued from previous page

Setting	Value for vidm-primary	Value for vidm-secondary-1	Value for vidm-secondary-2
FQDN	Enter the FQDN for vidm-primary	Enter the FQDN for vidm-secondary-1.	Enter the FQDN for vidm-secondary-2.
IP address	Enter the IP Address for vidm-primary.	Enter the IP Address for vidm-secondary-1.	Enter the IP Address for vidm-secondary-2.

e) For each node, click advanced configuration and click **Select Root Password**.

Select `xint-wsa-root` and click **Save**.

11. On the **Precheck** page, click **Run precheck**.
12. On the **Manual validations** page, select the **I took care of the manual steps above and am ready to proceed** check box and click **Run precheck**.
13. Review the validation report, remediate any errors, and click **Re-run precheck**.
14. Wait for all prechecks to complete with `Passed` messages and click **Next**.
15. On the **Summary** page, review the configuration details. To back up the deployment configuration, click **Export configuration**.
16. To start the deployment, click **Submit**.
The **Request details** page displays the progress of deployment.
17. Monitor the steps of the deployment graph until all stages become `Completed`.

Configure an Anti-Affinity Rule and a Virtual Machine Group for a Clustered Workspace ONE Access Instance

To protect the nodes in a clustered Workspace ONE Access instance from a host-level failure, configure an anti-affinity rule to run the virtual machines on different hosts in the default management vSphere cluster. You then configure a VM group to define the startup order to ensure that vSphere High Availability powers on the clustered Workspace ONE Access nodes in the correct order.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. In the Hosts and Clusters inventory, expand the management domain vCenter Server and data center.
3. Select the cluster and click the **Configure** tab.
4. Create the anti-affinity rule for the clustered Workspace ONE Access virtual machines.
 - a) Navigate to **Configuration > VM/Host rules** and click **Add**.
 - b) Configure the settings and click **OK**.

Setting	Value
Name	<management-domain-name>-anti-affinity-rule-wsa
Enable rule	Selected
Type	Separate Virtual Machines
Members	Click Add , select the clustered Workspace ONE Access nodes, and click OK . <ul style="list-style-type: none"> • vidm-primary_VM • vidm-secondary-1_VM

Table continued on next page

Continued from previous page

Setting	Value
	• vidm-secondary-2_VM

5. Create a virtual machine group for the clustered Workspace ONE Access nodes.
 - a) Navigate to **Configuration > VM/Host groups** and click **Add**.
 - b) Configure the settings and click **OK**.

Setting	Value
Name	Clustered Workspace ONE Access Appliances
Type	VM Group
Members	Click Add , select the clustered Workspace ONE Access nodes, and click OK . <ul style="list-style-type: none"> • vidm-primary_VM • vidm-secondary-1_VM • vidm-secondary-2_VM

Configure NTP on Workspace ONE Access

To keep NTP synchronized with the other SDDC components, configure NTP using the Workspace ONE Access appliance configuration interface.

1. In a web browser, log in to the Workspace ONE Access instance with the **admin** user by using the appliance configuration interface (https://<wsa_node_fqdn>:8443/cfg/login).
2. In the navigator pane, click **Time synchronization**.
3. Configure the settings and click **Save**.

Setting	Description
Time sync	NTP selected
NTP Server	Enter the FQDN of the NTP server.

4. If you deployed a cluster, repeat this procedure for the remaining clustered Workspace ONE Access nodes.

Configure the Domain and Domain Search Parameters on Workspace ONE Access

To enable name translation and resolution between the region-specific and the cross-region environments, configure the domain name and domain search parameters on Workspace ONE Access.

1. Log in to the cross-region Workspace ONE Access instance by using a Secure Shell (SSH) client.
2. Switch to the super user by running the `su` command.
3. Open the `/etc/resolv.conf` file in a text editor.

```
vi /etc/resolv.conf
```

4. Add entries for `Domain` and `search` to the end of the file and save the file. For example:


```
Domain rainpole.io
search rainpole.io sfo.rainpole.io
```

5. If you deployed a clustered Workspace ONE Access instance, repeat this procedure for the remaining nodes in the cluster.

Configure an Identity Source for Workspace ONE Access

To enable identity and access management in the SDDC, you integrate your Active Directory with Workspace ONE Access and configure attributes to synchronize users and groups.

1. In a web browser, log in to Workspace ONE Access by using the administration interface to the **System Domain** with **configadmin** user (https://<wsa_fqdn>/admin).
2. On the main navigation bar, click **Identity and access management**.
3. Click the **Directories** tab, and from the **Add directory** drop-down menu, select **Add Active Directory over LDAP/IWA**.
4. On the **Add directory** page, configure the following settings, click **Test connection** and click **Save and next**.

Setting	Value
Directory name	Enter a name for directory. For example, <code>sfo.rainpole.io</code> .
Active Directory over LDAP	Selected
Sync connector	Select the FQDN of <code>vidm-primary</code>
Do you want this connector to also perform authentication?	Yes
Directory search attribute	<code>SAMAccountName</code>
This Directory requires all connections to use STARTTLS (Optional)	If you want to secure communication between Workspace ONE Access and Active Directory select this option and paste the Root CA certificate in the SSL Certificate box.
Base DN	Enter the Base Distinguished Name from which to start user searches. For example, <code>cn=Users,dc=sfo,dc=rainpole,dc=io</code> .
Bind DN	Enter the DN for the user to connect to Active Directory. For example, <code>cn=svc-wsa-ad,ou=Service Accounts,dc=sfo,dc=rainpole,dc=io</code> .
Bind user password	Enter the password for the Bind user. For example: <code>svc-wsa-ad_password</code> .

5. On the **Select the domains** page, review the domain name and click **Next**.
6. On the **Map user attributes** page, review the attribute mappings and click **Next**.
7. On the **Select the groups (users) you want to sync** page, enter the distinguished name for the folder containing your groups (For example `OU=Security Groups,DC=sfo,DC=rainpole,DC=io`) and click **Select**.

8. For each **Group DN** you want to include, select the group to use by Workspace ONE Access for each of the roles, and click **Save** then **Next**.

Product	Role Assigned via Group
Workspace ONE Access	Super Admin
	Directory Admin
	ReadOnly Admin
VMware Aria Suite Lifecycle	VCF Role
	Content Admin
	Content Developers

9. On the **Select the Users you would like to sync** page, enter the distinguished name for the folder containing your users (e.g. OU=Users, DC=sfo, DC=rainpole, DC=io) and click **Next**.
10. On the **Review** page, click **Edit**, from the **Sync frequency** drop-down menu, select **Every 15 minutes**, and click **Save**.
11. To initialize the directory import, click **Sync directory**.

Add the Clustered Workspace ONE Access Cluster Nodes as Identity Provider Connectors

To provide high availability for the identity and access management services of a clustered Workspace ONE Access instance, you add the cluster nodes as directory connectors.

This procedure is only applicable if you deployed a clustered Workspace ONE Access instance. It does not apply to a standard Workspace ONE Access instance.

- In a web browser, log in to the clustered Workspace ONE Access instance by using the administration interface to the **System Domain** with **configadmin** user (https://<wsa_cluster_fqdn>/admin).
- On the main navigation bar, click **Identity and access management**.
- Click the **Identity Providers** tab.
- Click the **WorkspaceIDP__1** identity provider.
- On the **WorkspaceIDP__1 details** page, under **Connector(s)** from the **Add a connector** drop-down menu, select `vidm-secondary-1_VM`, configure the settings, and click **Add connector**.

Setting	Value
Connector	<code>vidm-secondary-1_VM</code>
Bind to AD	Checked
Bind user password	<code>svc-wsa-ad_password</code>

- Repeat this step for the `vidm-secondary-2_VM` connector.
- In the **IdP Hostname** text box, enter the FQDN of the NSX load balancer virtual server for Workspace ONE Access cluster.
- Click **Save**.

Assign Roles to Active Directory Groups for Workspace ONE Access

Workspace ONE Access uses role-based access control to manage delegation of roles. You assign the **Super Admin**, **Directory Admin** and **ReadOnly** roles to Active Directory groups to manage access to Workspace ONE Access.

You assign the following administrator roles to the corresponding user groups.

Workspace ONE Access Role	Example Active Directory Group Name
Super Admin	wsa-admins
Directory Admin	wsa-directory-admin
ReadOnly Admin	wsa-read-only

1. In a web browser, log in to Workspace ONE Access by using the administration interface to the System Domain with **configadmin** user (https://<wsa_fqdn>/admin).
2. On the main navigation bar, click **Roles**.
3. Assign Workspace ONE Access roles to Active Directory groups.
 - a) Select the **Super Admin** role and click **Assign**.
 - b) In the **Users / User Groups** search box, enter the name of the Active Directory group you want to assign the role to, select the group, and click **Save**.
 - c) Repeat this step to configure the **Directory Admin** and the **ReadOnly Admin** roles.

Assign Roles to Active Directory Groups for VMware Aria Suite Lifecycle

To enable identity and access management for VMware Aria Suite Lifecycle, you integrate the component with the clustered Workspace ONE Access instance.

You assign the following administrative roles to corresponding Active Directory groups.

VMware Aria Suite Lifecycle Role	Example Active Directory Group Name
VCF Role	vrslcm-admins
Content Release Manager	vrslcm-release-manager
Content Developer	vrslcm-content-developer

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Identity and Tenant Management**.
3. In the navigation pane, click **User management** and click **Add user / group**.
4. On the **Select users / groups** page, in the search box, enter the name of the group you want to assign the role to, select the Active Directory group, and click **Next**.
5. On the **Select roles** page, select the **VCF Role** role, and click **Next**.
6. On the **Summary** page, click **Submit**.
7. Repeat this procedure to assign roles to the **Content Release Manager** and **Content Developer** user groups.

Working with NSX Federation in VMware Cloud Foundation

With NSX Federation, you can federate NSX environments across VMware Cloud Foundation (VCF) instances. You can manage federated NSX environments with a single pane of glass, create gateways and segments that span VMware Cloud Foundation instances, and configure and enforce firewall rules consistently across instances.

IMPORTANT

If you plan to deploy VMware Aria Suite components, you must deploy Application Virtual Networks before you configure NSX Federation. See [Deploying Application Virtual Networks in VMware Cloud Foundation](#).

NSX Federation is supported between VCF and non-VCF deployments. If you choose to federate NSX between VCF and non-VCF deployments, you are responsible for the deployment and lifecycle of the NSX Global Managers, as well as maintaining version interoperability between VCF-owned NSX Local Managers, non-VCF NSX Local Managers, and the NSX Global Manager.

NSX Federation Key Concepts

NSX Federation introduces some new terms and concepts in VMware Cloud Foundation (VCF).

NSX Federation Systems: Global Manager and Local Manager

An NSX Federation environment within VMware Cloud Foundation includes two types of management systems.

Global Manager: a system similar to NSX Manager that federates multiple Local Managers.

Local Manager: an NSX Manager system in charge of network and security services for a VMware Cloud Foundation instance.

NSX Federation Span: Local and Cross-Instance

When you create a networking object from Global Manager, it can span one or more VMware Cloud Foundation instances.

Local: the object spans only one instance.

Cross-instance: the object spans more than one instance. You do not directly configure the span of a segment. A segment has the same span as the gateway it is attached to.

NSX Federation Tunnel Endpoints

In an NSX Federation environment, there are two types of tunnel endpoints.

Tunnel End Point (TEP): the IP address of a transport node (Edge node or Host) used for Geneve encapsulation within an instance.

Remote Tunnel End Points (RTEP): the IP address of a transport node (Edge node only) used for Geneve encapsulation across instances.

NSX Federation Tier Gateways

An NSX Federation in VMware Cloud Foundation environment includes three types of tier-1 gateways.

Type	Description	Managed By	Scope
standalone tier-1 gateway	Configured in the Local Manager and used for services such as the Load Balancer.	Local Manager	Single VMware Cloud Foundation instance
local-instance tier-1 gateway	Configured in the Global Manager at a single location, this is a global tier-1 gateway used for segments that exist within a single	Global Manager	Single VMware Cloud Foundation instance

Table continued on next page

Continued from previous page

Type	Description	Managed By	Scope
	VMware Cloud Foundation Instance.		
cross-instance tier-1 gateway	Configured in the Global Manager, this is a global Tier-1 gateway used for segments that exist across multiple VMware Cloud instances.	Global Manager	Multiple VMware Cloud Foundation instance

Configuring NSX Federation in VMware Cloud Foundation

With NSX Federation, you can federate the management domain NSX or a VI workload domain NSX across VMware Cloud Foundation (VCF) instances.

See [VMware Configuration Maximums](#) for your version of NSX for information about the maximum number of supported federated NSX Managers and other NSX federation maximums.

NOTE

VI workload domains that share an NSX Manager are considered a single location.

Some tasks described in this section are to be performed on the first NSX instance while others need to be performed on each NSX instance that is being federated. See the table below for more information.

NSX Instance	Tasks to be Performed
First Instance	<ol style="list-style-type: none"> 1. Create Global Manager Clusters for 2. Replacing Global Manager Cluster Certificates in You can skip this step if you are using self-signed certificates. 3. Prepare Local Manager for NSX Federation in 4. Enable NSX Federation in 5. Prepare for Stretching Segments between VMware Cloud Foundation Instances: <ol style="list-style-type: none"> a. Create and Configure Cross-Instance Tier-1 Gateway b. Connect Cross-Instance Segments to Cross-Instance Tier-1 Gateway
Enable high availability for NSX Federation Control Plane on one additional instance	<ol style="list-style-type: none"> 1. Create Global Manager Clusters for 2. Replacing Global Manager Cluster Certificates in You can skip this step if you are using self-signed certificates. 3. Set Standby Global Manager

Table continued on next page

Continued from previous page

NSX Instance	Tasks to be Performed
Each additional instance	<ol style="list-style-type: none"> 1. Prepare Local Manager for NSX Federation in 2. Add Location to Global Manager 3. Stretching Segments between VMware Cloud Foundation Instances: <ol style="list-style-type: none"> a. Delete Existing Tier-0 Gateways in Additional Instances b. Connect Additional VMware Cloud Foundation Instances to Cross-Instance Tier-0 Gateway c. Connect Local Tier-1 Gateway to Cross-Instance Tier-0 Gateway d. Add Additional Instance as Locations to the Cross-Instance Tier-1 Gateway

Create Global Manager Clusters for VMware Cloud Foundation

An NSX Federation environment contains an active and a standby Global Manager cluster and one or more Local Manager clusters. The standby Global Manager appliance provides high availability and disaster recovery.

- [Set Active Global Manager](#)
- [Set Standby Global Manager](#)

Deploy Global Manager Nodes

You deploy three Global Manager nodes in the VMware Cloud Foundation management domain.

1. Download the NSX OVF file from the VMware download portal.
2. In a web browser, log in to vCenter Server at https://vcenter_server_fqdn/ui.
3. Select the default cluster in the management domain.
4. Right-click and select **Deploy OVF template**.
5. Select **Local file**, click **Upload files**, and navigate to the OVA file.
6. Click **Next**.
7. Enter a name and a location for the NSX Manager VM, and click **Next**.

The name you enter appears in the vSphere and vCenter Server inventory.

8. Select the compute resource on which to deploy the NSX Manager appliance page and click **Next**.
9. Review and verify the OVF template details and click **Next**.
10. Accept the license agreement and click **Next**.
11. Specify the deployment configuration size and click **Next**.

The Description panel on the right side of the wizard shows the details of selected configuration. You can also refer to [VMware Configuration Maximums](#) to ensure that you choose the correct size for the scale or your environment.

12. Specify storage for the configuration and disk files.
 - Select the virtual disk format.
 - Select the VM storage policy.
 - Specify the datastore to store the NSX Manager appliance files.
 - Click **Next**.

NOTE

The virtual disk format is determined by the selected VM storage policy when using a vSAN datastore.

13. Select the management network as the destination network and click **Next**.

The following steps are all located in the Customize Template section of the Deploy OVF Template wizard.

14. In the Application section, enter the system root, CLI admin, and audit passwords for the NSX Manager. The root and admin credentials are mandatory fields.

Your passwords must comply with the password strength restrictions.

- At least 12 characters
- At least one lower-case letter
- At least one upper-case letter
- At least one digit
- At least one special character
- At least five different characters

15. In the Optional parameters section, leave the password fields blank.
16. In the Network Properties section, enter the hostname of the NSX Manager.
17. For Rolename, select the NSX Global Manager role.
18. Enter the default gateway, management network IPv4, and management network netmask.
19. In the DNS section, enter the DNS Server list and Domain Search list.
20. In the Services Configuration section, enter the NTP Server list and enable SSH.
21. Verify that all your custom OVF template specification is accurate and click **Finish** to initiate the deployment.

The deployment might take 7-8 minutes.

22. After the deployment is complete, power on the Global Manager node.

Right-click the Global Manager VM and, from the **Actions** menu, select **Power > Power on**.

23. In a web browser, log in to Global Manager at https://gm_node1_fqdn/.
24. Accept the end-user license agreement and click **Continue**.
25. Join the VMware Customer Experience Program and click **Save**.
26. Repeat steps 4 - 22 to deploy two additional Global Manager nodes.

Join Global Manager Nodes to Form a Cluster

Join the three Global Manager nodes you deployed in the VMware Cloud Foundation management domain to form a cluster.

1. SSH into the first NSX Global Manager node using the `admin` user account.
2. Run the following command to retrieve the Global Manager cluster ID.

```
get cluster config | find Id:
```

3. Copy the output of the command and save it.
4. Run the following command to retrieve the thumbprint of the Global Manager API certificate.

```
get certificate api thumbprint
```

5. Copy the output of the command and save it.
6. Log in to the second Global Manager node and run the following command to join this node to the cluster:

```
join first_node_IP cluster-id cluster_ID username admin password nsx_admin_password
thumbprint api_thumbprint
```

where *cluster_ID* is the value from step 3 and *certificate_thumbprint* is the value from step 5.

A warning message displays: Data on this node will be lost. Are you sure? (yes/no).

7. Enter **yes** to confirm.

The joining and cluster stabilizing process might take from 10 to 15 minutes.

8. Run **get cluster status** to view the status.

Verify that the status for every cluster service group is UP before making any other cluster changes.

9. Repeat steps 6-8 to join the third node to the cluster.

10. Verify the cluster status on the web interface.

- a) Log in to the Global Manager web interface and select **Configuration > Global Manager Appliances**.
- b) Verify that the **Cluster status** is green that the cluster node is **Available**.

Create Anti-Affinity Rule for Global Manager Cluster in VMware Cloud Foundation

Create an anti-affinity rule to ensure that the Global Manager nodes run on different ESXi hosts. If an ESXi host is unavailable, the Global Manager nodes on the other hosts continue to provide support for the NSX management and control planes.

1. In a web browser, log in to the management domain or VI workload domain vCenter Server at https://vcenter_server_fqdn/ui.
2. Select **Menu > Hosts and Clusters**.
3. In the inventory, expand **vCenter Server > Datacenter**.
4. Select the Global Manager cluster and click the **Configure** tab.
5. Select **VM/Host rules** and click **Add**.
6. Enter the rule details.

Option	Description
Name	Type a name for the rule.
Enable rule	Select this option.
Type	Select Separate Virtual Machines .
Members	Click Add , select the three Global Manager nodes, and click OK .

7. Click **OK** in the Create VM/Host rule dialog box.

Assign a Virtual IP Address to Global Manager Cluster

To provide fault tolerance and high availability to Global Manager nodes, assign a virtual IP address (VIP) to the Global Manager cluster in VMware Cloud Foundation.

1. In a web browser, log in to a Global Manager node at https://gm_node_1-fqdn/.
2. Click **System** and then select **Global Manager Appliances**.
3. Click **Set Virtual IP** and enter the VIP address for the cluster. Ensure that VIP is part of the same subnet as the other management nodes.
4. Click **Save**.
5. Verify that the VIP is working correctly.

From a browser, log in to the Global Manager using the virtual IP address assigned to the cluster at `https://gm_vip_fqdn/`.

Prepare Local Manager for NSX Federation in VMware Cloud Foundation

To prepare for NSX Federation, you create an IP pool in the Local Manager. The Global Manager assigns IP addresses from this pool to the Edge nodes for remote tunnel end point (RTEP) interfaces. You also set the global fabric MTU to match the end-to-end MTU between instances.

1. In a web browser, log in to Local Manager cluster for the management domain or VI workload domain at `https://lm_vip_fqdn/`.
2. On the main navigation bar, click **Networking**.
3. Create an IP pool for RTEP in Local Manager
 - a) In the navigation pane, select **IP Address Pools** and click **Add IP address pool**.
 - b) Enter a name.
 - c) Under Subnets, click **Set**.
 - d) In the Set Subnets dialog box, click **Add subnet > IP Ranges**.
 - e) Configure the settings and click **Add**.
 - f) Click **Add** and then click **Apply**.
 - g) Click **Save**.
4. Configure MTU for RTEP.
 - a) On the main navigation bar, click **System**.
 - b) Select **Fabric > Settings**.
 - c) Under **Global Fabric Settings**, Click **Edit** for Remote Tunnel Endpoint.
 - d) Enter **9000** in MTU and click **Save**.

Enable NSX Federation in VMware Cloud Foundation

To enable NSX Federation in VMware Cloud Foundation, set the Global Manager as active and add the existing NSX Manager in the management domain or VI workload domain as a location to the Global Manager.

Set Active Global Manager

Activate the Global Manager.

1. In a web browser, log in to Global Manager cluster for the management or VI workload domain at `https://gm_vip_fqdn/`.
2. Click **System** and then select **Location Manager**.
3. Click **Make Active** and enter a name for the active Global Manager.
4. Click **Save**.

Add Location to Global Manager

Add the NSX Manager in the management domain or VI workload domain as a location to the Global Manager. This NSX Manager is now referred to as the Local Manager. You then import segments, tier-0 gateways, and tier-1 gateways from the Local Manager to the Global Manager.

1. Obtain the certificate thumbprint of the NSX Local Manager cluster.
 - a) Enable SSH on one of the NSX Manager VMs.
 - b) From the vCenter UI, open the web console of one of the NSX Managers and login to the Admin user.
 - c) Run the command `start service ssh` to enable SSH on the NSX Manager.
 - d) Use a Secure Shell (SSH) client and log in to the same NSX Manager with the Admin user.

- e) Run the command `get certificate cluster thumbprint` to retrieve the Local Manager cluster VIP thumbprint.

```
sfo-m01-nsxt01c> get certificate cluster thumbprint
```

```
b88c4e052fe61309915527511e7f1b25970286a51cf1dd68ea881daba1ed0a9f
```

- f) Save the thumbprint.
 g) Run the `stop service ssh` command to deactivate SSH on the NSX Manager.
2. Add NSX Manager as a location to the Global Manager.
- a) Log in to Global Manager at https://active_gm_vip_fqdn/.
 b) Select **System** > **Location Manager** and click **Add On-Prem Location**.
 c) In the Add New Location dialog box, enter the location details.

Option	Description
Location Name	Enter a name for the location.
FQDN/IP	Enter the FQDN or IP address of the NSX Manager cluster VIP. Do not enter an individual NSX Manager FQDN or IP.
Username and Password	Provide the admin user's credentials for the NSX Manager at the location.
SHA-256 Thumbprint	Add the thumbprint you retrieved in step 1.
Check Compatibility	Click Check Compatibility to ensure that the location can be added. This checks that the NSX version is compatible.

- d) Click **Save**
3. Configure networking on the Local Manager nodes.
- a) On the Location Manager page, in the Locations section, click **Networking** under the location you are adding then click **Configure**.
 b) On the Configure Edge Nodes for Stretch Networking page, click **Select All**
 c) In the Remote Tunnel Endpoint Configuration pane enter the following details.

Option	Value
Host Switch	nsxDefaultHostSwitch
Teaming Policy Name	Select Use Default .
RTEP VLAN	Enter the VLAN for the host.
IP Pool for all Nodes	Select the IP pool.

- d) Click **Save**.
4. Import the Local Manager configuration to the Global Manager.
- a) Select the Global Manager context from the drop down menu.

NOTE

You may need to refresh your browser or logout and log in to the Global Manager to see the drop down menu.

- b) On the System tab, select the Location Manager pane.
 c) Under **Locations**, click **Import**.

This option may take 15 minutes or longer to appear.

- d) Verify that you have a recent backup and click **Proceed to import**.
- e) In the Preparing for import dialog box, click **Next** and then click **Import**.
Wait for a confirmation that the import is successful.

Local Manager objects imported into the Global Manager are owned by the Global Manager and appear in the Local Manager with a GM icon. You can modify these objects only from the Global Manager.

5. Repeat these steps for each Local Manager cluster.

Stretch Segments between VMware Cloud Foundation Instances

Each NSX Manager instance to be federated has a tier-0 gateway, tier-1 gateway, and two segments created during NSX Edge deployment and Application Virtual Network (AVN) creation. One of these segments is for local instance use and the other is for cross-instance use. Both segments are initially connected to the same tier-1 gateway. When NSX Manager instances are federated, you create an additional tier-1 gateway for cross-instance use and migrate the cross-instance segment from the original tier-1 gateway to the new tier-1 gateway. The new tier-1 gateway has locations for both instances enabled on it. This allows you to manage the ingress and egress routing for cross-instance segments when you move them between VMware Cloud Foundation instances independently of local instance segments whose ingress and egress remain unaffected.

NOTE

Cross-instance segments cannot have overlapping IP addresses/ranges.

Create and Configure Cross-Instance Tier-1 Gateway

You create a new tier-1 gateway in one of the VMware Cloud Foundation instances. You then extend this gateway to the other federated instances.

1. In a web browser, log in to Global Manager for the management or VI workload domain at https://gm_vip_fqdn/.
2. On the main navigation bar, click **Networking**.
3. In the navigation pane, select **Tier-1 gateways**.
4. Specify the gateway details.

Setting	Specified Value
Tier-1 Gateway Name	Enter a name for the new tier-1 gateway.
Linked Tier-0 Gateway	Enter the global tier-0 gateway.
Edges Pool Allocation Size	Select Routing .
Enable Edge Clusters for Services or Custom span	Select Enabled .
Fail Over	Select Non Preemptive .
Enable Standby Relocation	Select Enabled .
Edge Cluster	Select the Edge cluster.
Mode	Select Primary

5. Click **Save**.
6. Click **Yes** to continue the configuration of the tier-1 gateway.
7. Configure route advertisement for the tier-1 gateway.
 - a) Expand the **Route advertisement** section of the tier-1 gateway.
 - b) Enable all available sources, click **Save**, and click **Close editing**.

Connect Cross-Instance Segments to Cross-Instance Tier-1 Gateway

You connect the cross-instance segments in the first instance to the cross-instance tier-1 gateway you created.

1. In a web browser, log in to Global Manager cluster at https://gm_vip_fqdn/.
2. On the NSX Manager main navigation bar, click **Networking**.
3. In the navigation pane, select **Segments**.
4. On the Segments tab, click the vertical eclipses for the *cross-instance_nsx_segment* and click **Edit**.
5. Change the Connected Gateway from *instance_tier1* to *cross-instance_tier1*, click **Save**, and then click **Close editing**.

Delete Existing Tier-0 Gateways in Additional Instances

Since you will use the cross-instance tier-0 gateway for upstream connections, you delete the local tier-0 gateway from each additional VCF instance.

1. In a web browser, log in to Global Manager cluster at https://active_gm_vip_fqdn/.
2. On the NSX Manager main navigation bar, click **Networking**.
3. Disconnect the tier-1 gateway for the NSX Local Manager.
 - a) In the navigation pane, select Tier-1 Gateways.
 - b) On the Tier-1 Gateways tab, click the vertical eclipses for the *additional_instance_tier1_gateway* and click **Edit**.
 - c) Under Linked Tier-0 gateway, click the X to disconnect the *additional_instance_tier0_gateway*, click **Save**, and click **Close editing**.

CAUTION

At this point any segments connected to *additional_instance_tier1_gateway* will be unreachable until you have finished connecting the additional instance to the cross-instance tier-0 infrastructure.

4. In the navigation pane, select Tier-0 Gateways.
5. On the Tier-0 Gateway page, click the vertical eclipses for the *additional_instance_tier0_gateway* and click **Delete**.
6. Click **Delete**.

Connect Additional VMware Cloud Foundation Instances to Cross-Instance Tier-0 Gateway

You turn the standard tier-0 gateway into a cross-instance tier-0 gateway by connecting additional VMware Cloud Foundation instances to it. You configure uplink interfaces, BGP, and route redistribution for the additional instances.

1. In a web browser, log in to Global Manager cluster at https://active_gm_vip_fqdn/.
2. Add the additional instance as a location on the tier-0 gateway.
 - a) On the NSX Manager main navigation bar, click **Networking**.
 - b) In the navigation pane, select **Tier-0 Gateways**.
 - c) On the Tier-0 Gateway page, click the vertical eclipses for the *cross-instance_tier0_gateway* and click **Edit**.
 - d) Click **Add Location** and enter the required information.

Setting	Value
Location	Select the location name of the instance being added.
Edge Cluster	Select the Edge cluster name of the instance being added.

- e) Click **Save**.
3. Set interfaces for the instance on the tier-0 gateway.
 - a) Expand **Interfaces** and click **Set**.
 - b) Click **Add interface**.
 - c) Enter a name for the interface and select the instance location.
 - d) Set the type to **External** and enter the IP address for the interface.
 - e) Select the segment that the interface is connected to and the Edge node corresponding to the instance.
 - f) Set the MTU to 9000.
 - g) Repeat these steps to add three additional interfaces.
4. Configure BGP neighbors.
 - a) Expand BGP and under BGP Neighbors, click **Set**.
You can enable BFD if the network supports it and is configured for BFD.
 - a) Click **Add BGP neighbor**
 - b) Enter the IP address for the neighbor and select the instance location.
 - c) Enter the remote AS and source addresses for the neighbor.
 - d) Click **Timers & Password** and set the **Hold Down Time** to 12 and **Keep Alive Time** to 4.
 - e) Enter the BGP neighbor password, click **Save**, and then click **Close**.
 - f) Repeat these steps to add another BGP neighbor.
5. Configure Route Re-Distribution
 - a) Expand Route Re-Distribution and next to the location you are adding, click **Set**.
 - b) In the Set Route Re-distribution dialog box, click **Add Route-Redistribution**.
 - c) Enter *default* as name and, under Route re-distribution, click **Set**.
 - d) In the Set route redistribution dialog box, select all listed sources and click **Apply**.
 - e) Click **Add** to finish editing the default route redistribution and click **Apply**.
 - f) Click **Save**
6. Click **Close editing**.

Connect Local Tier-1 Gateway to Cross-Instance Tier-0 Gateway

You connect the local tier-1 gateway at each VCF instance to the cross-instance tier-0 gateway.

1. In a web browser, log in to Global Manager cluster at https://active_gm_vip_fqdn/.
2. On the NSX Manager main navigation bar, click **Networking**.
3. In the navigation pane, select **Tier-1 gateways**.
4. On the Tier-1 Gateway page, click the vertical ellipses menu for the *this_instance_tier1_gateway* and click **Edit**.
5. Change the Connected Gateway to *cross_instance_tier0_gateway*.
6. In the Location change dialog box, click **Yes**.
7. Under Locations, delete all locations except the location of the instance you are working with.
8. Click **Save** and click **Close Editing**.

Add Additional Instance as Locations to the Cross-Instance Tier-1 Gateway

Add each additional instance as a location on the cross-instance Tier-1 gateway to enable cross-instance workloads.

1. In a web browser, log in to Global Manager cluster at `https://active_gm_vip_fqdn/`.
2. On the NSX Manager main navigation bar, click **Networking**.
3. In the navigation pane, select **Tier-1 Gateways**.
4. On the Tier-1 Gateway page, click the vertical eclipses for the `cross-instance_tier1` gateway and click **Edit**.
5. Click **Add Location** and enter the following values.

Setting	Value
Location	Select the location of this instance
Edge Cluster	Select the NSX Edge cluster of the this instance
Mode	Set to Secondary .

6. Click **Save** and click **Close Editing**.

Set Standby Global Manager

You provide high availability of the active Global Manager by configuring the Global Manager in the additional instance as standby to the active cluster. In case of failure of the cluster in first instance, you can use the cluster in additional instance to provide the networking capabilities.

Create the standby Global Manager cluster. See [Create Global Manager Clusters for VMware Cloud Foundation](#).

1. Obtain the certificate thumbprint of the Standby Global Manager cluster.
 - a) Enable SSH on one of the NSX Manager VMs.
 - b) From the vCenter UI, open the web console of one of the NSX Managers and login to the Admin user.
 - c) Run the command `start service ssh` to enable SSH on the NSX Manager.
 - d) Use a Secure Shell (SSH) client and log in to the same NSX Manager with the Admin user.
 - e) Run the command `get certificate cluster thumbprint` to retrieve the Global Manager cluster thumbprint.


```
sfo-m01-nsxt01c> get certificate cluster thumbprint

b88c4e052fe61309915527511e7f1b25970286a51cf1dd68ea881daba1ed0a9f
```
 - f) Save the thumbprint.
 - g) Run the `stop service ssh` command to deactivate SSH on the NSX Manager.
2. Add additional Global Manager instance
 - a) Log in to the Active Global Manager at `https://active_gm_vip_fqdn/`.
 - b) On the main navigation bar, Select **System > Location Manager**.
 - c) Click **Add Standby**.
 - d) Enter the location name, FQDN, username and password, and the SHA-256 thumbprint you had retrieved earlier.
 - e) Click **Check Compatibility** and click **Save**.

Replacing Global Manager Cluster Certificates in VMware Cloud Foundation

To replace certificates for the Global Manager cluster, you import root and intermediate CA-signed certificates as appropriate and replace the Global Manager default certificates with the imported certificates using API calls.

Import a CA-Signed Certificate to the Global Manager Cluster

Import the root/leaf or machine certificate and intermediate certificate as appropriate to the first Global Manager node.

Generate root and intermediate CA-signed certificates.

1. In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
2. Import the root CA certificate.
 - a) On the main navigation bar, **System > Certificates**.
 - b) Click **Import > Import CA certificate**.
 - c) In the Import CA Certificate dialog box, enter a name for the root CA certificate.
 - d) For **Certificate Contents**, select the root CA certificate you created in step 2c and click **Import**.
3. Import certificates for the Global Manager nodes and the load balanced virtual server address.
 - a) Click **Import > Import certificate**.
 - b) In the **Name** field, enter *gm_vip_fqdn*.
 - c) In the Certificate Contents, browse to the previously created certificate file with the extension *chain.pem* and select the file.
 - d) In the **Private Key**, browse to the previously created private key with the extension *.key*, select the file, and click **Import**.

Replace the Certificate for the First Global Manager Node

Replace the default certificate of the first Global Manager node to establish a trusted connection with the management components in the SDDC. You use APIs for this procedure.

1. In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
2. Retrieve the certificate ID.
 - a) On the main navigation bar, click **System > Certificates**.
 - b) Copy the certificate ID value and save it.
3. Log in to the host that has access to your data center.
4. Replace the default certificate on the first Global Manager node with the CA-signed certificate.
 - a) Start the Postman application in your web browser and log in.
 - b) On the **Authorization** tab, enter the following settings.

Setting	Value
Type	Select Basic Auth .
User name	Enter <i>admin</i> .
Password	Enter <i>nsx_admin_password</i> .

- c) Click **Update request**.
- d) On the Headers tab, add a key as follows.

Setting	Value
Key	Content-Type
Key Value	application/xml

- e) In the request pane at the top, send the following HTTP request.

Setting	Value
HTTP request method	Select POST .
URL	Enter <code>https://gm_node1_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>

After the Global Manager sends a response, a 200 OK status is displayed on the Body tab.

5. Restart the first Global Manager node.
 - a) Log in to vCenter Server.
 - b) In the inventory expand **vCenter Server** › **Datacenter** › **Cluster**.
 - c) Right-click the node and select **Actions** › **Power** › **Restart guest OS**.

Replace Certificates and Virtual IP for the Remaining Global Manager Nodes

Replace the default certificates on the remaining Global Manager nodes.

Table 188: URLs for Replacing the Global Manager Node Certificates

NSX Manager Node	POST URL for Certificate Replacement
<code>gm_node2_fqdn</code>	<code>https://gm_node2_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>
<code>gm_node3_fqdn</code>	<code>https://gm_node3_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=gm_fqdn_certificate_ID</code>
<code>gm_vip_fqdn</code>	<code>https://gm_vip_fqdn/api/v1/cluster/api-certificate?action=set_cluster_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>

1. In a web browser, log in to the active Global Manager at `https://gm_vip_fqdn/`.
2. Log in to the host that has access to your data center.
3. Replace the default certificate for the second Global Manager node with the CA-signed certificate by using the first Global Manager node as a source.
 - a) Start the Postman application in your web browser and log in.
 - b) On the **Authorization** tab, configure the following settings.

Setting	Value
Type	Select Basic Auth.
User name	Enter <code>admin</code> .
Password	Enter the <code>nsx_admin_password</code> .

- a) Click **Update request**.
- b) On the **Headers** tab, enter the header details.

Setting	Value to Select
Key	Content-Type
Key Value	application/xml

c) In the request pane at the top, send the URL query.

Setting	Value
HTTP request method	Select POST.
URL	Enter <code>https://gm_node2_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=firstinstance_gm_vip_certificate_ID</code>

After the NSX Manager appliance responds, the Body tab displays a 200 OK status.

4. To upload the CA-signed certificate on the third Global Manager node, repeat steps 2 to step 4 with appropriate values.
5. Restart the second and third Global Manager nodes.
 - a) Log in to vCenter Server.
 - b) In the inventory expand **vCenter Server > Datacenter > Cluster**
 - c) Right-click the second and third Global Manager nodes and click **Actions > Power > Restart guest OS**.
6. Verify the status of each Global Manager node.
 - a) In a web browser, log in to the first Global Manager node at `https://gm_node1_fqdn/`.
 - b) For each node, navigate to **System > Global Manager Appliances > View Details** and confirm that the status is **REPO_SYNC = SUCCESS**.
7. Assign a certificate to the Global Manager cluster.
 - a) Start the Postman application in your web browser and log in.
 - b) On the **Authorization** tab, configure the following settings.

Setting	Value
Type	Select Basic Auth .
User name	Enter <code>admin</code> .
Password	Enter <code>nsx_admin_password</code> .

- c) Click **Update request**.
- d) On the Headers tab, add a key as follows.

Setting	Value
Key	Content-Type
Key Value	application/xml

e) In the request pane at the top, send the URL query.

Setting	Value
HTTP request method	Select POST .
URL	Enter <code>https://gm_vip_fqdn/api/v1/cluster/api-certificate?action=set_cluster_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>

After the NSX Global Manager sends a response, a 200 OK status is displayed on the Body tab.

Update Local Manager Certificate Thumbprint in Global Manager Cluster

After you rotate the Local Manager certificates using SDDC Manager, you obtain the new certificate thumbprint to update it in the Global Manager cluster.

1. In a web browser, log in to Global Manager at `https://nsx_gm_vip_fqdn/`.
2. Obtain certificate thumbprint.
 - a) Log in to a vCenter Server by using a Secure Shell (SSH) client.
 - b) Run the **shell** command to switch to the bash shell.
 - c) Run the command to retrieve the SHA-256 thumbprint of the virtual IP for the NSX Manager cluster certificate.

```
echo -n | openssl s_client -connect nsx_lm_vip_fqdn:443 2>/dev/null | openssl x509 -noout -fingerprint -sha256
```

- d) Save the thumbprint value.
3. Update the Local Manager certificate thumbprint in the Global Manager.
 - a) On the main navigation bar, click **System**.
 - b) In the navigation pane, select **Location Manager**.
 - c) Under **Locations**, select the Local Manager instance, and click **Actions**.
 - d) Click **Edit Settings** and update NSX Local Manager Certificate Thumbprint.
 - e) Click **Check Compatibility** and click **Save**.
 - f) Wait for the Sync Status to display success and verify that all Local Manager nodes appear.
 4. Under Locations, update the Local Manager certificate thumbprint for all the instances.

Password Management for NSX Global Manager Cluster in VMware Cloud Foundation

You can manage NSX Global Manager user accounts using the NSX appliance's CLI. Resetting the password for any of the local users on one node automatically resets the password for the other NSX Managers in the cluster. The synchronization of the password can take a few minutes.

See [Manage Local User's Password or Name Using the CLI](#).

Backup and Restore of NSX Global Manager Cluster in VMware Cloud Foundation

Regular backups of the NSX Global Manager components ensures that you can keep your environment operational if a data loss or failure occurs.

The Global Manager cluster stores the configured state of the segments. If the Global Manager appliances become unavailable, the network traffic in the data plane is intact but you can make no configuration changes.

Configure NSX Global Manager Cluster Backups

Configure an SFTP server to store backup files. After a backup file server is configured, you can start a backup at any time, or schedule recurring backups.

1. In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
2. Select **System > Backup & Restore**.
3. On the Backup tab, click **Edit**.
4. Enter the IP address or FQDN of the backup file server.
5. Change the default port if necessary. The default port is 22.
6. The protocol text box is already filled in. SFTP is the only supported protocol.
7. In the **Directory Path** text box, enter the absolute directory path where the backups will be stored.
8. Enter the user name and password required to log in to the backup file server.

The first time you configure a file server, you must provide a password. Subsequently, if you reconfigure the file server, and the server IP or FQDN, port, and user name are the same, you do not need to enter the password again.

9. Leave the **SSH Fingerprint** blank and accept the fingerprint provided by the server after you click Save in a later step.
10. Enter a passphrase.

NOTE

You will need this passphrase to restore a backup. If you forget the passphrase, you cannot restore any backups.

11. Click Edit under the Schedule label.

You can schedule recurring backups or trigger backups for configuration changes.

1. Click the Recurring Backup toggle.
2. Click Weekly and set the days and time of the backup, or click Interval and set the interval between backups.
3. Enabling the **Detect NSX configuration change** option will trigger an unscheduled full configuration backup when it detects any runtime or non-configuration related changes, or any change in user configuration. For Global Manager, this setting triggers backup if any changes in the database are detected, such as the addition or removal of a Local Manager or Tier-0 gateway or DFW policy.
4. You can specify a time interval for detecting database configuration changes. The valid range is 5 minutes to 1,440 minutes (24 hours). This option can potentially generate a large number of backups. Use it with caution.
5. Click **Save**.

After you configure a backup file server, you can click **Backup Now** to manually start a backup at any time. Automatic backups run as scheduled. You see a progress bar of your in-progress backup.

Restore an NSX Global Manager Cluster Backup

Restoring a backup restores the state of the network at the time of the backup. In addition, the configurations maintained by Global Manager appliances are also restored.

- Verify that you have the login credentials for the backup file server.
- Verify that you have the SSH fingerprint of the backup file server. Only SHA256 hashed ECDSA (256 bit) host key is accepted as a fingerprint.
- Verify that you have the passphrase of the backup file.

Do not change the configuration of the NSX Global Manager cluster while the restore process is in progress.

1. If any nodes in the appliance cluster that you are restoring are online, power them off.
2. Install one new appliance node on which to restore the backup.
 - If the backup listing for the backup you are restoring contains an IP address, you must deploy the new Global Manager node with the same IP address. Do not configure the node to publish its FQDN.
 - If the backup listing for the backup you are restoring contains an FQDN, you must configure the new appliance node with this FQDN and publish the FQDN. Only lowercase FQDN is supported for backup and restore.
3. In a web browser, log in to Global Manager at `https://gm_vip_fqdn/`.
4. Make the Global Manager active. You can restore a backup only on an active Global Manager.
 - a) On the main navigation bar, click **System**.
 - b) In the navigation pane, select **Location Manager**.
 - c) On the Location Manager page, click **Make Active**, enter a name for the Global Manager, and click **Save**.
5. On the main navigation bar, click **System** > **Backup & Restore** and then click **Edit**.
6. Enter the IP address or FQDN of the backup file server.
7. Change the default port if necessary. The default port is 22.
8. To log in to the server, enter the user name and password.
9. In the **Destination Directory** text box, enter the absolute directory path where the backups are stored.
10. Enter the passphrase that was used to encrypt the backup data.
11. Leave the **SSH Fingerprint** blank and accept the fingerprint provided by the server after you click Save in a later step.
12. Select a backup and click **Restore**.
13. The restore process prompts you to take action, if necessary, as it progresses.
14. After the restored manager node is up and functional, deploy additional nodes to form a NSX Global Manager cluster.

Managing Installation and Upgrade Bundles in VMware Cloud Foundation

Installation bundles include the software required to deploy VMware Aria Suite Lifecycle, NSX Advanced Load Balancer, or a VI workload domain in VMware Cloud Foundation. Upgrade bundles enable you to perform updates on SDDC Manager, vCenter Server, ESXi, and NSX. Bundles can be downloaded and applied manually or scheduled within your maintenance window.

Bundle Types

VMware Cloud Foundation includes two types of bundles.

Install Bundles

An install bundle includes software binaries to install VI workload domains, VMware Aria Suite Lifecycle, and NSX Advanced Load Balancer.

VMware Cloud Foundation includes the following install bundles:

- A VI workload domain install bundle is used to deploy later versions of the software components instead of the versions in your original VMware Cloud Foundation installation. It includes software binaries for vCenter Server and NSX.
- The VMware Aria Suite Lifecycle install bundle is used for deploying VMware Aria Suite Lifecycle.

NOTE

For other products in the VMware Aria Suite, you can download the installation binaries directly from VMware Aria Suite Lifecycle. See "Configure Product Binaries" in the *VMware Aria Suite Lifecycle Installation, Upgrade, and Management Guide* for the version of [VMware Aria Suite Lifecycle](#) listed in the VMware Cloud Foundation BOM.

- A bundle for installing NSX Advanced Load Balancer in a workload domain.

Update or Upgrade Bundles

An update or upgrade bundle includes software binaries to update the appropriate VMware Cloud Foundation software components. In most cases, an upgrade bundle must be applied to the management domain before it can be applied to VI workload domains.

For information about update/upgrade bundles, see the *VMware Cloud Foundation Lifecycle Management Guide*.

NOTE

See [Public URL list for SDDC Manager](#) for information about the URLs that must be accessible to download bundles.

Downloading Install Bundles for VMware Cloud Foundation

You can download install bundles directly from the SDDC Manager UI if Lifecycle Management is connected to an online or offline depot. If SDDC Manager does not have direct internet connectivity, you can either use a proxy server to access the online depot, set up an offline depot, or download the bundles manually using the Bundle Transfer Utility. For information about downloading upgrade bundles, see the *VMware Cloud Foundation Lifecycle Management Guide*.

Connect SDDC Manager to a Software Depot for Downloading Bundles

SDDC Manager can connect to a software depot to download software bundles, compatibility data, and more.

To connect to the online depot, SDDC Manager must be able to connect to the internet, either directly or through a proxy server.

To connect to an offline depot, you must first configure it. See [KB 312168](#) for information about the requirements and process for creating an offline depot. To download bundles to an offline depot, see "Download Bundles to an Offline Depot" in the *VMware Cloud Foundation Lifecycle Management Guide*.

SDDC Manager supports two types of software depots:


- Online depot
- Offline depot

You can only connect SDDC Manager to one type of depot. If SDDC Manager is connected to an online depot and you configure a connection to an offline depot, the online depot connection is disabled and deleted.

1. In the navigation pane, click **Administration** › **Depot Settings**.

Online Depots


VMware Depot

 Depot connection not set up. Authenticate your Customer Connect Account to set it up.

[AUTHENTICATE](#)

Offline Depot

Offline Depot

 Depot connection not set up.

[SET UP](#)

2. Connect SDDC Manager to an online depot or an offline depot.

Depot Type	Configuration Steps
Online	<ol style="list-style-type: none"> 1. Click Authenticate for the VMware Depot. 2. Type your Broadcom Support Portal user name and password. 3. Click Authenticate
Offline	<ol style="list-style-type: none"> 1. Click Set Up for the Offline Depot. 2. Enter the following information for the offline depot: <ul style="list-style-type: none"> – FQDN or IP address – Port number – User name – Password 3. Click Set Up.

SDDC Manager attempts to connect to the depot. If the connection is successful, SDDC Manager starts looking for available bundles. To view available bundles, click **Lifecycle Management > Bundle Management** and then click the **Bundles** tab. It may take some time for all available bundles to appear.

Download an Install Bundle from SDDC Manager

You can download install bundles from the SDDC Manager UI.

Connect SDDC Manager to an online or offline depot. See [Connect SDDC Manager to a Software Depot for Downloading Bundles](#).

If you do not have direct internet access, configure a proxy server or use the Bundle Transfer Utility for offline bundle downloads.

- [Configure a Proxy Server for Downloading Bundles](#)
- [Download an Install Bundle Using the Bundle Transfer Utility](#)

1. In the navigation pane, click **Lifecycle Management** › **Bundle Management**.
2. Click the **Bundles** tab.
All available bundles are displayed. If you recently set up your depot connection, it may take some time for all available bundles to appear. Install bundles display an **Install Only Bundle** label.
3. For the bundle you want to download, do one of the following:
 - For an immediate download, click **Download Now**.
The bundle download begins right away.
 - To schedule a download, click **Schedule Download**.
Select the date and time for the bundle download and click **Schedule**.
4. Navigate to **Lifecycle Management** › **Bundle Management** › **Download History** to see the downloaded bundles.

Configure a Proxy Server for Downloading Bundles

If you do not have direct internet access, you can configure a proxy server to download bundles. VMware Cloud Foundation 5.2 and later support proxy servers with authentication.

1. In the navigation pane, click **Administration** › **Proxy Settings**.
2. Click **Set Up Proxy**.
3. Toggle the **Enable Proxy** setting to the on position.
4. Select **HTTP** or **HTTPS**.
5. Enter the proxy server FQDN or IP address and port number.
6. If your proxy server requires authentication, toggle the **Authentication** setting to the on position and enter the user name and password.
7. Click **Save**.

You can now download install bundles as described in [Download an Install Bundle from SDDC Manager](#).

Download an Install Bundle Using the Bundle Transfer Utility

If you do not have direct internet connectivity to your SDDC Manager instance, you can use the Bundle Transfer Utility to manually download install bundles from the VMware depot to your local computer and then upload them to SDDC Manager.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles.
- The computer must have Java 8 or later.
- A Windows or Linux computer with access to the SDDC Manager appliance for uploading the bundles.
- To upload the manifest file from a Windows computer, you must have OpenSSL installed and configured.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.

NOTE

The Bundle Transfer Utility is the only supported method for downloading bundles. Do not use third-party tools or other methods to download bundles.

This procedure describes the process for downloading install bundles using the Bundle Transfer Utility. For information about downloading update/upgrade bundles, see the *VMware Cloud Foundation Lifecycle Management Guide*. If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN or IP address and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

- Download the most recent version of the Bundle Transfer Utility on a computer with internet access.
 - Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
 - Click the version of VMware Cloud Foundation to which you are upgrading.
 - Click **Drivers & Tools**.
 - Click the download icon for the Bundle Transfer Utility.
- Extract `lcm-tools-prod.tar.gz`.
- Navigate to the `lcm-tools-prod/bin/` and confirm that you have execute permission on all folders.
- Copy the Bundle Transfer Utility to a computer with access to the SDDC Manager appliance and then copy the bundle transfer utility to the SDDC Manager appliance.
 - SSH in to the SDDC Manager appliance using the `vcf` user account.
 - Enter `su` to switch to the root user.
 - Create the `lcm-tools` directory.

```
mkdir /opt/vmware/vcf/lcm/lcm-tools
```

NOTE

If the `/opt/vmware/vcf/lcm/lcm-tools` directory already exists with an older version of the Bundle Transfer Utility, you need to delete contents of the existing directory before proceeding.

- Copy the Bundle Transfer Utility file (`lcm-tools-prod.tar.gz`) that you downloaded in step 1 to the `/opt/vmware/vcf/lcm/lcm-tools` directory.
- Extract the contents of `lcm-tools-prod.tar.gz`.

```
tar -xvf lcm-tools-prod.tar.gz
```


- f) Set the permissions for the `lcm-tools` directory.

```
cd /opt/vmware/vcf/lcm/
chown vcf_lcm:vcf -R lcm-tools
chmod 750 -R lcm-tools
```

5. On the computer with internet access, download the manifest file. This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username
```

6. Copy the manifest file and `lcm-tools-prod` directory to a computer with access to the SDDC Manager appliance.

7. Upload the manifest file to the SDDC Manager appliance.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory Manifest-Downloaded-Directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

8. Download install bundles from the computer with internet access.

- a) From the `/lcm-tools/bin` folder where you downloaded the utility download install bundles by entering the following command.

Windows:

```
lcm-bundle-transfer-util.bat -download --outputDirectory absolute-path-output-dir
-depotUser depotUser -p vcfVersion --imageType
```

INSTALL

Linux:

```
./lcm-bundle-transfer-util -download --outputDirectory absolute-path-output-dir
-depotUser depotUser -p vcfVersion --imageType
```

INSTALL

For example:

```
./lcm-bundle-transfer-util -download --outputDirectory /root/downloadedBundles
-depotUser ffroyven@vmware.com -p 4.4.0.0 --imageType INSTALL
```

where

<code>absolute-path-output-dir</code>	Path to the directory where the bundle files are to be downloaded. This directory folder must have 777 permissions. If you do not specify the download directory, bundles are downloaded to the default directory with 777 permissions.
<code>depotUser</code>	Broadcom Support Portal user name. You are prompted to enter the password. If there are any special characters in the password, specify the password within single quotes.

Table continued on next page

Continued from previous page

-p	Filter the bundles for a specific version of VMware Cloud Foundation. The value is based on x.x.x.x format.
----	---

After you enter your Broadcom Support Portal password, the utility asks `Do you want to download vRealize bundles?`. Enter `Y` or `N`. The utility displays a list of the available install bundles for the specified version of VMware Cloud Foundation.

9. Specify the bundles to download.

Enter one of the following options:

- `all`
- A specific bundle name or a comma-separated list of bundle names to download specific bundles. For example: `bundle-52610, bundle-52990`.

10. Copy the entire output directory to a computer with access to the SDDC Manager appliance, and then copy it to the SDDC Manager appliance.

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

Example command SDDC Manager appliance

```
scp -pr /root/downloadedBundles vcf@SDDC_MANAGER_IP:/nfs/vmware/vcf/nfs-mount/
```

The `scp` command in the example above copies the output directory (`downloadedBundles`) to the `/nfs/vmware/vcf/nfs-mount/` directory on the SDDC Manager appliance.

11. Upload the directory to the SDDC Manager appliance internal LCM repository.

- a) SSH in to the SDDC Manager appliance using the `vcf` user account.
- b) Navigate to `/opt/vmware/vcf/lcm/lcm-tools/bin`.
- c) Run the following command:

```
./lcm-bundle-transfer-util -upload -bundleDirectory absolute-path-bundle-dir
```

Replace *absolute-path-bundle-dir* with the path to the location where you copied the output directory. For example: `/nfs/vmware/vcf/nfs-mount/downloadedBundles`.

The utility uploads the bundles and displays upload status for each bundle. When the uploads complete, the bundles are available in the SDDC Manager UI. Navigate to **Lifecycle Management** > **Bundle Management** > **Download History** to see the downloaded bundles.

View Bundle Download History

The SDDC Manager UI allows you to view the bundle download history. From here you can see what bundles are already downloaded and available.

1. In the navigation pane, click **Lifecycle Management** > **Bundle Management**.
2. Click the **Download History** tab.

Stretching vSAN Clusters in VMware Cloud Foundation

You can stretch a vSAN cluster (ESA or OSA) in a workload domain across two availability zones within a region. Both availability zones must contain an equal number of hosts to ensure failover in case any of the availability zones goes down.

The default management vSphere cluster must be stretched before a VI workload domain cluster can be stretched. This ensures that the NSX control plane and management VMs (vCenter, NSX management cluster, and SDDC Manager) remain accessible if the stretched cluster in the primary availability zone goes down.

NOTE

You cannot stretch a cluster in the following conditions:

- The cluster is a vSAN Max cluster.
- The cluster has a vSAN remote datastore mounted on it.
- The cluster shares a vSAN Storage Policy with any other clusters.
- The cluster includes DPU-backed hosts.

Some use cases for stretching a cluster are described below.

- **Planned maintenance**
You can perform a planned maintenance on an availability zone without any downtime and then migrate the applications after the maintenance is completed.
- **Automated recovery**
Stretching a cluster automatically initiates VM restart and recovery, and has a low recovery time for the majority of unplanned failures.
- **Disaster avoidance**
With a stretched cluster, you can prevent service outages before an impending disaster.

About Availability Zones and Regions

This section describes an availability zone and region as used for stretch clusters.

Availability Zones

An availability zone is a collection of infrastructure components. Each availability zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Each zone should have independent power, cooling, network, and security.

Additionally, these zones should be physically separate so that disasters affect only one zone. The physical distance between availability zones is short enough to offer low, single-digit latency (less than 5 ms) and large bandwidth (10 Gbps) between the zones.

Availability zones can either be two distinct data centers in a metro distance, or two safety or fire sectors (data halls) in the same large-scale data center.

Regions

Regions are in two distinct locations - for example, region A can be in San Francisco and region B in Los Angeles (LAX). Distance between regions can be larger than the distance between availability zones. The latency between regions must be less than 150 ms.

Stretched Cluster Requirements

In an environment with multiple availability zones, Layer 2 networks must be stretched between the availability zones by the physical infrastructure. You also must provide a Layer 3 gateway that is highly available between availability zones. The method for stretching these Layer 2 networks and providing a highly available Layer 3 gateway is vendor-specific.

VLANs and Subnets for Multiple Available Zones

This section displays a sample configuration for an environment with multiple availability zones. The VM management, Uplink 01, Uplink 02, and Edge overlay networks in each availability zone must be stretched to facilitate failover of the NSX Edge appliances between availability zones. The Layer 3 gateway for the management and Edge overlay networks must be highly available across the availability zones.

NOTE

If VLAN is stretched between AZ1 and AZ2, the Layer 3 network must also be stretched between the two AZs.

Table 189: Stretched Cluster Subnet Requirements

Function	Availability Zone 1	Availability Zone 2	HA Layer 3 Gateway	Recommended MTU
VM Management VLAN	✓	✓	✓	1500
Management VLAN (AZ1)	✓	X	✓	1500
vMotion VLAN	✓	X	✓	9000
vSAN VLAN (AZ1)	✓	X	✓	9000
NSX Host Overlay VLAN	✓	X	✓	9000
NSX Edge Uplink01 VLAN	✓	✓	X	9000
NSX Edge Uplink02 VLAN	✓	✓	X	9000
NSX Edge Overlay VLAN	✓	✓	✓	9000
Management VLAN (AZ2)	X	✓	✓	1500
vMotion VLAN (AZ2)	X	✓	✓	9000
vSAN VLAN (AZ2)	X	✓	✓	9000
NSX Host Overlay VLAN (AZ2)	X	✓	✓	9000

Networking for Multiple Availability Zones

There are specific physical data center network requirements for a topology with multiple availability zones. For information about the vSAN witness appliance requirements, see [vSAN Witness Design](#) for in the *VMware Cloud Foundation Design Guide*.

Table 190: Physical Network Requirements for Multiple Availability Zone

Component	Requirement
MTU	<p>VLANS which are stretched between availability zones must meet the same requirements as the VLANS for intra-zone connection including MTU. MTU value must be consistent end-to-end including components on the inter-zone networking path. Set MTU values as follows.</p> <ul style="list-style-type: none"> • MTU for all VLANS and Switch Virtual Interfaces (vMotion, Geneve, and Storage) to jumbo frames. • Management MTU to 1500. • Geneve overlay requires a minimum MTU of 1600.
Layer 3 gateway availability	For VLANS that are are stretched between available zones, configure data center provided method to failover the Layer

Table continued on next page

Continued from previous page

Component	Requirement
	3 gateway between availability zones. For example, VRRP or HSRP.
DHCP availability	For VLANs that are stretched between availability zones, provide high availability for the DHCP server so that a failover operation of a single availability zone will not impact DHCP availability.
BGP routing	Each availability zone data center must have its own Autonomous System Number (ASN).
Ingress and egress traffic	<ul style="list-style-type: none"> • For VLANs that are stretched between availability zones, traffic flows in and out of a single zone. Local egress is not supported. • For VLANs that are not stretched between availability zones, traffic flows in and out of the zone where the VLAN is located. • For NSX virtual network segments that are stretched between regions, traffic flows in and out of a single availability zone. Local egress is not supported.
Latency	<p>vSphere</p> <ul style="list-style-type: none"> • Less than 150 ms latency RTT for vCenter Server connectivity. • Less than 150 ms latency RTT for vMotion connectivity. • Less than 5 ms latency RTT for VSAN hosts connectivity. <p>vSAN</p> <ul style="list-style-type: none"> • Less than 200 ms latency RTT for up to 10 hosts per site. • Less than 100 ms latency RTT for 11-15 hosts per site. <p>NSX Managers</p> <ul style="list-style-type: none"> • Less than 10 ms latency RTT between NSX Managers • Less than 150 ms latency RTT between NSX Managers and transport nodes.

Deploy and Configure vSAN Witness Host

Each vSAN stretched cluster requires a witness host deployed in a vSAN witness zone, which must be different from the location of both availability zones.

You deploy the vSAN witness host using an appliance instead of using a dedicated physical ESXi host as a witness host. The witness host does not run virtual machines and must run the same version of ESXi as the ESXi hosts in the stretched cluster. It must also meet latency and Round Trip Time (RTT) requirements.

There are separate vSAN witness appliances for vSAN OSA and vSAN ESA. You must deploy the witness appliance that matches the cluster type that you are stretching.

See the Physical Network Requirements for Multiple Availability Zone table within [Stretched Cluster Requirements](#).

Deploy vSAN Witness Host

You deploy the vSAN witness host for a stretched cluster at a site which is isolated from the existing availability zones to prevent propagation of failure or outage in the data center.

Download the VMware vSAN Witness Appliance .ova file from the [Broadcom Support Portal](#).

- For stretching a vSAN OSA cluster download the appliance for vSAN OSA.
- For stretching a vSAN ESA cluster download the appliance for vSAN ESA.

For more information, see [vSAN Witness Design for VMware Cloud Foundation](#).

1. In a web browser, log in to vCenter Server at `https://vcenter_server_fqdn/ui`.
2. Select **Menu** > **Hosts and Clusters**.
3. In the inventory panel, expand **vCenter Server** > **Datacenter**.
4. Right-click the cluster and select **Deploy OVF template**.
5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, browse to the location of the vSAN witness host OVA file, and click **Next**.
6. On the **Select a name and folder** page, enter a name for the virtual machine and click **Next**.
7. On the **Select a compute resource** page, click **Next**.
8. On the **Review details** page, review the settings and click **Next**.
9. On the **License agreements** page, accept the license agreement and click **Next**.
10. On the **Configuration** page, select **Medium** and click **Next**.
11. On the **Select storage** page, select a datastore and click **Next**.
12. On the **Select networks** page, select a portgroup for the witness and management network, and click **Next**.
13. On the **Customize template** page, enter the root password for the witness and click **Next**.
14. On the **Ready to complete** page, click **Finish** and wait for the process to complete.
15. Power on the vSAN witness host.
 - a) In the inventory panel, navigate to **vCenter Server** > **Datacenter** > **Cluster**.
 - b) Right-click the vSAN witness host and from the **Actions** menu, select **Power** > **Power on**.

Register vSAN Witness Host

Before you can configure the vSAN Witness Host, you must register it with vCenter Server.

1. Use the vSphere Client to log in to the vCenter Server containing the cluster that you want to stretch.
2. In the vSphere Client, navigate to the data center.
3. Right-click the data center and select **Add Host**.

IMPORTANT

You must add the vSAN Witness Host to the datacenter. Do not add it to a folder.

4. On the **Name and location** page, enter the Fully Qualified Domain Name (FQDN) of the vSAN Witness Host and click **Next**.

NOTE

Do not use the IP address.

5. On the **Connection settings** page, enter administrator credentials and click **Next**.
6. On the **Host summary** page, review the summary of the host details and click **Next**.
7. On the **Host lifecycle** page, the check box **Manage host with an image** is selected by default.
 - If you want to manage the host with an image, leave the check box selected and click **Next**.
 - If you do not want to manage the host with an image, deselect the check box and click **Next**.
8. If you manage the host with an image, on the **Image** page, set up the desired image and click **Next**.

- On the **Assign license** page, assign an existing license and click **Next**.

NOTE

Do not create a new license.

- Review the summary and click **Finish**.

Configure NTP on the Witness Host

To prevent time synchronization issues, configure the NTP service on the vSAN witness host.

- In the inventory panel of the vCenter Server Client, select **vCenter Server > Datacenter**.
- Select the vSAN witness host and click the **Configure** tab.
- Configure the NTP client on the vSAN witness host.
 - In the **System** section, click **Time configuration** and click the **Edit** button.
 - Select **Use Network Time Protocol (enable NTP client)**.
 - Configure the following settings and click **OK**.

Setting	Value
NTP Servers	NTP server address
Start NTP Service	Selected
NTP Service Startup Policy	Start and stop with host

Configure the VMkernel Adapters on the vSAN Witness Host

To enable vSAN data network communication between the availability zones, configure the witness network on the vSAN witness host.

- In the inventory panel of the vCenter Server Client, select **vCenter Server > Datacenter**.
- Select the vSAN witness host and click the **Configure** tab.
- Remove the dedicated witness traffic VMkernel adapter on the vSAN Witness host.
 - In the **Networking** section, click **VMkernel adapters**.
 - Select the kernel adapter **vmk1** with `secondaryPg` as **Network label** and click **Remove**.
 - On the **Remove VMkernel adapter** dialog box, click **Remove**.
- Remove the virtual machine network port group on the vSAN witness host.
 - In the left pane, select **Networking > Virtual switches**.
 - Expand the **Standard switch: secondary switch** section.
 - Click the vertical ellipsis and from the drop-down menu, select **Remove**.
 - On the **Remove standard switch** dialog box, click **Yes**.
 - Expand the **Standard switch: vSwitch0** section.
 - In the **VM Network** pane, click the vertical ellipsis and from the drop-down menu, select **Remove**.
 - On the **Remove port group** dialog box, click **Yes**.
- Enable witness traffic on the VMkernel adapter for the management network of the vSAN witness host.
 - On the **VMkernel adapters** page, select the **vmk0** adapter and click **Edit**.
 - In the **vmk0 - edit settings** dialog box, click **Port properties**, select the **vSAN** check box, and click **OK**.

Stretch a vSAN Cluster in VMware Cloud Foundation

You can stretch a vSAN cluster (ESA or OSA) in the management domain or VI workload domain using a JSON specification and the VMware Cloud Foundation API.

- Verify that vCenter Server is operational.
- Verify that you have completed the Planning and Preparation Workbook with the management domain or VI workload domain deployment option included.
- Verify that your environment meets the requirements listed in the Prerequisite Checklist sheet in the Planning and Preparation Workbook.
- Create a network pool for availability zone 2.
- Commission hosts for availability zone 2. See [Commission Hosts](#).
- Ensure that you have enough hosts such that there is an equal number of hosts on each availability zone. This is to ensure that there are sufficient resources in case an availability zone goes down completely.
- Deploy and configure a vSAN witness host. See [Deploy and Configure vSAN Witness Host](#).
- If you are stretching a cluster in a VI workload domain, the default management vSphere cluster must have been stretched.

NOTE

You cannot stretch a cluster in the following cases:

- The cluster is a vSAN Max cluster.
- The cluster has a vSAN remote datastore mounted on it.
- The cluster shares a vSAN Storage Policy with any other clusters.
- The cluster includes DPU-backed hosts.

When you stretch a cluster, VMware Cloud Foundation modifies the site disaster tolerance setting for storage policy associated with datastore of that cluster from **None - standard cluster** to **Site mirroring - stretched cluster**. This affects all VMs using default datastore policy in that cluster. If you do not want to change the site disaster tolerance setting for specific VMs, apply a different storage policy to those VMs before stretching the cluster.

1. Create a JSON specification in a text editor.

The following example is for an environment with a single vSphere Distributed Switch. If you have multiple vSphere Distributed Switches, see the [VMware Cloud Foundation API Reference Guide](#) for details about creating a JSON specification.

NOTE

The ESXi hosts that you are adding to availability zone 2 must use the same vmnic to vSphere Distributed Switch mapping as the existing hosts in availability zone 1.

```
{
  "clusterStretchSpec": {
    "hostSpecs": [
      {
        "hostname": "sfo02-w01-esx01.sfo.rainpole.io",
        "hostNetworkSpec": {
          "networkProfileName": "sfo-w01-az2-nsx-np01",
          "vmNics": [
            {
              "id": "vmnic0",
```



```
    "uplink": "uplink1",
    "vdsName": "sfo-w01-cl01-vds01"
  },
  {
    "id": "vmnic1",
    "uplink": "uplink2",
    "vdsName": "sfo-w01-cl01-vds01"
  }
]
},
"id": "<ESXi host 1 ID>",
"licenseKey": "<license key>"
},
{
  "hostname": "sfo02-w01-esx02.sfo.rainpole.io",
  "hostNetworkSpec": {
    "networkProfileName": "sfo-w01-az2-nsx-np01",
    "vmNics": [
      {
        "id": "vmnic0",
        "uplink": "uplink1",
        "vdsName": "sfo-w01-cl01-vds01"
      },
      {
        "id": "vmnic1",
        "uplink": "uplink2",
        "vdsName": "sfo-w01-cl01-vds01"
      }
    ]
  }
},
"id": "<ESXi host 2 ID>",
"licenseKey": "<license key>"
```

```
},
{
  "hostname": "sfo02-w01-esx03.sfo.rainpole.io",
  "hostNetworkSpec": {
    "networkProfileName": "sfo-w01-az2-nsx-np01",
    "vmNics": [
      {
        "id": "vmnic0",
        "uplink": "uplink1",
        "vdsName": "sfo-w01-cl01-vds01"
      },
      {
        "id": "vmnic1",
        "uplink": "uplink2",
        "vdsName": "sfo-w01-cl01-vds01"
      }
    ]
  },
  "id": "<ESXi host 3 ID>",
  "licenseKey": "<license key>"
},
{
  "hostname": "sfo02-w01-esx04.sfo.rainpole.io",
  "hostNetworkSpec": {
    "networkProfileName": "sfo-w01-az2-nsx-np01",
    "vmNics": [
      {
        "id": "vmnic0",
        "uplink": "uplink1",
        "vdsName": "sfo-w01-cl01-vds01"
      },
      {
```

```
    "id": "vmnic1",
    "uplink": "uplink2",
    "vdsName": "sfo-w01-cl01-vds01"
  }
]
},
"id": "<ESXi host 4 ID>",
"licenseKey": "<license key>"
}
],
"isEdgeClusterConfiguredForMultiAZ": <true, if the cluster hosts an NSX Edge
cluster; false, if the cluster does not host an NSX Edge cluster>,
"networkSpec": {
  "networkProfiles": [
    {
      "isDefault": false,
      "name": "sfo-w01-az2-nsx-np01",
      "nsxtHostSwitchConfigs": [
        {
          "ipAddressPoolName": "sfo-w01-az2-host-ip-pool01",
          "uplinkProfileName": "sfo-w01-az2-host-uplink-profile01",
          "vdsName": "sfo-w01-cl01-vds01",
          "vdsUplinkToNsxUplink": [
            {
              "nsxUplinkName": "uplink-1",
              "vdsUplinkName": "uplink1"
            },
            {
              "nsxUplinkName": "uplink-2",
              "vdsUplinkName": "uplink2"
            }
          ]
        }
      ]
    }
  ]
}
```

```
]
}
],
"nsxClusterSpec": {
  "ipAddressPoolsSpec": [
    {
      "description": "WLD01 AZ2 Host TEP Pool",
      "name": "sfo-w01-az2-host-ip-pool01",
      "subnets": [
        {
          "cidr": "172.16.44.0/24",
          "gateway": "172.16.44.253",
          "ipAddressPoolRanges": [
            {
              "end": "172.16.44.200",
              "start": "172.16.44.10"
            }
          ]
        }
      ]
    }
  ],
  "uplinkProfiles": [
    {
      "name": "sfo-w01-az2-host-uplink-profile01",
      "teamings": [
        {
          "activeUplinks": [
            "uplink-1",
            "uplink-2"
          ],
          "name": "DEFAULT",
```

```

    "policy": "LOADBALANCE_SRCID",
    "standByUplinks": []
  }
],
"transportVlan": 1644
}
]
}
},
"witnessSpec": {
  "fqdn": "sfo-w01-cl01-vsw01.sfo.rainpole.io",
  "vsanCidr": "172.17.11.0/24",
  "vsanIp": "172.17.11.219"
},
"witnessTrafficSharedWithVsanTraffic": false
}
}

```

NOTE

Replace the example values in the JSON file with the correct values for your environment.

2. In the navigation pane, click **Developer Center > API Explorer**.
3. Retrieve and replace the unique IDs for each ESXi host in the JSON specification.
 - a) Expand the **APIs for managing hosts** section, and expand **GET /v1/hosts**.
 - b) In the **Status** text box, enter `UNASSIGNED_USEABLE` and click **Execute**.
 - c) In the **Response** section, click **PageOfHost**, copy the `id` element of each host, and replace the respective value in the JSON specification.

ESXi Host	Value
ESXi Host 1	<i>ESXi host 1 ID</i>
ESXi Host 2	<i>ESXi host 2 ID</i>
ESXi Host 3	<i>ESXi host 3 ID</i>
ESXi Host 4	<i>ESXi host 4 ID</i>

4. Replace the license key value in JSON specification with valid keys.
5. Retrieve the unique ID for the management cluster.
 - a) Expand the **APIs for managing clusters** section, and expand **GET /v1/cluster**.
 - b) Click **Execute**.

- c) In the **Response** section, click **PageOfCluster**, copy the `id` element of the management cluster.
You will need the cluster ID for subsequent steps.

6. Validate the JSON specification file.

- a) Expand the **APIs for managing clusters** section and expand **POST /v1/clusters/{id}/validations**.
b) In the **Value** text box, enter the unique ID for the management cluster that you retrieved in step 5.
c) In the **clusterUpdateSpec** text box, type

```
{
  "clusterUpdateSpec":
}
```

- d) Paste the JSON specification.

For example:

```
{
  "clusterUpdateSpec": {
    "clusterStretchSpec": {
      "hostSpecs": [
        {
          "hostname": "sfo02-w01-esx01.sfo.rainpole.io",
          "hostNetworkSpec": {
            "networkProfileName": "sfo-w01-az2-nsx-np01",
            "vmNics": [
              {
                "id": "vmnic0",
                "uplink": "uplink1",
                "vdsName": "sfo-w01-cl01-vds01"
              },
              {
                "id": "vmnic1",
                "uplink": "uplink2",
                "vdsName": "sfo-w01-cl01-vds01"
              }
            ]
          },
          "id": "<ESXi host 1 ID>",
          "licenseKey": "<license key>"
        }
      ]
    }
  }
}
```

```
},
{
  "hostname": "sfo02-w01-esx02.sfo.rainpole.io",
  "hostNetworkSpec": {
    "networkProfileName": "sfo-w01-az2-nsx-np01",
    "vmNics": [
      {
        "id": "vmnic0",
        "uplink": "uplink1",
        "vdsName": "sfo-w01-cl01-vds01"
      },
      {
        "id": "vmnic1",
        "uplink": "uplink2",
        "vdsName": "sfo-w01-cl01-vds01"
      }
    ]
  },
  "id": "<ESXi host 2 ID>",
  "licenseKey": "<license key>"
},
{
  "hostname": "sfo02-w01-esx03.sfo.rainpole.io",
  "hostNetworkSpec": {
    "networkProfileName": "sfo-w01-az2-nsx-np01",
    "vmNics": [
      {
        "id": "vmnic0",
        "uplink": "uplink1",
        "vdsName": "sfo-w01-cl01-vds01"
      },
      {
```

```

        "id": "vmnic1",
        "uplink": "uplink2",
        "vdsName": "sfo-w01-cl01-vds01"
    }
]
},
"id": "<ESXi host 3 ID>",
"licenseKey": "<license key>"
},
{
    "hostname": "sfo02-w01-esx04.sfo.rainpole.io",
    "hostNetworkSpec": {
        "networkProfileName": "sfo-w01-az2-nsx-np01",
        "vmNics": [
            {
                "id": "vmnic0",
                "uplink": "uplink1",
                "vdsName": "sfo-w01-cl01-vds01"
            },
            {
                "id": "vmnic1",
                "uplink": "uplink2",
                "vdsName": "sfo-w01-cl01-vds01"
            }
        ]
    },
    "id": "<ESXi host 4 ID>",
    "licenseKey": "<license key>"
}
],
"isEdgeClusterConfiguredForMultiAZ": <true, if the cluster hosts an NSX Edge
cluster; false, if the cluster does not host an NSX Edge cluster>,
"networkSpec": {

```



```
"networkProfiles": [
  {
    "isDefault": false,
    "name": "sfo-w01-az2-nsx-np01",
    "nsxtHostSwitchConfigs": [
      {
        "ipAddressPoolName": "sfo-w01-az2-host-ip-pool01",
        "uplinkProfileName": "sfo-w01-az2-host-uplink-profile01",
        "vdsName": "sfo-w01-cl01-vds01",
        "vdsUplinkToNsxUplink": [
          {
            "nsxUplinkName": "uplink-1",
            "vdsUplinkName": "uplink1"
          },
          {
            "nsxUplinkName": "uplink-2",
            "vdsUplinkName": "uplink2"
          }
        ]
      }
    ]
  }
],
"nsxClusterSpec": {
  "ipAddressPoolsSpec": [
    {
      "description": "WLD01 AZ2 Host TEP Pool",
      "name": "sfo-w01-az2-host-ip-pool01",
      "subnets": [
        {
          "cidr": "172.16.44.0/24",
          "gateway": "172.16.44.253",
```

```
    "ipAddressPoolRanges": [
      {
        "end": "172.16.44.200",
        "start": "172.16.44.10"
      }
    ]
  }
]
}
],
"uplinkProfiles": [
  {
    "name": "sfo-w01-az2-host-uplink-profile01",
    "teamings": [
      {
        "activeUplinks": [
          "uplink-1",
          "uplink-2"
        ],
        "name": "DEFAULT",
        "policy": "LOADBALANCE_SRCID",
        "standByUplinks": []
      }
    ],
    "transportVlan": 1644
  }
]
},
"witnessSpec": {
  "fqdn": "sfo-w01-cl01-vsw01.sfo.rainpole.io",
  "vsanCidr": "172.17.11.0/24",
```

```

    "vsanIp": "172.17.11.219"
  },
  "witnessTrafficSharedWithVsanTraffic": false
}
}

```

NOTE

Replace the example values in the JSON file with the correct values for your environment.

- e) Click **Execute**.
 - f) In the confirmation dialog box, click **Continue**.
 - g) In the **Response** section, expand the **result** section and verify that the response is `SUCCEEDED`.
7. Stretch the cluster with the JSON specification.
- a) Expand the **APIs for managing clusters** section and expand **PATCH /v1/clusters/{id}**.
 - b) Paste the unique `ID` of the management cluster in the **Value** text-box.
 - c) In the **clusterUpdateSpec** text box, paste the JSON specification.

For example:

```

{
  "clusterStretchSpec": {
    "hostSpecs": [
      {
        "hostname": "sfo02-w01-esx01.sfo.rainpole.io",
        "hostNetworkSpec": {
          "networkProfileName": "sfo-w01-az2-nsx-np01",
          "vmNics": [
            {
              "id": "vmnic0",
              "uplink": "uplink1",
              "vdsName": "sfo-w01-cl01-vds01"
            },
            {
              "id": "vmnic1",
              "uplink": "uplink2",
              "vdsName": "sfo-w01-cl01-vds01"
            }
          ]
        }
      }
    ]
  }
}

```

```
    },
    "id": "<ESXi host 1 ID>",
    "licenseKey": "<license key>"
  },
  {
    "hostname": "sfo02-w01-esx02.sfo.rainpole.io",
    "hostNetworkSpec": {
      "networkProfileName": "sfo-w01-az2-nsx-np01",
      "vmNics": [
        {
          "id": "vmnic0",
          "uplink": "uplink1",
          "vdsName": "sfo-w01-cl01-vds01"
        },
        {
          "id": "vmnic1",
          "uplink": "uplink2",
          "vdsName": "sfo-w01-cl01-vds01"
        }
      ]
    },
    "id": "<ESXi host 2 ID>",
    "licenseKey": "<license key>"
  },
  {
    "hostname": "sfo02-w01-esx03.sfo.rainpole.io",
    "hostNetworkSpec": {
      "networkProfileName": "sfo-w01-az2-nsx-np01",
      "vmNics": [
        {
          "id": "vmnic0",
          "uplink": "uplink1",
```

```
    "vdsName": "sfo-w01-cl01-vds01"
  },
  {
    "id": "vmnic1",
    "uplink": "uplink2",
    "vdsName": "sfo-w01-cl01-vds01"
  }
]
},
"id": "<ESXi host 3 ID>",
"licenseKey": "<license key>"
},
{
  "hostname": "sfo02-w01-esx04.sfo.rainpole.io",
  "hostNetworkSpec": {
    "networkProfileName": "sfo-w01-az2-nsx-np01",
    "vmNics": [
      {
        "id": "vmnic0",
        "uplink": "uplink1",
        "vdsName": "sfo-w01-cl01-vds01"
      },
      {
        "id": "vmnic1",
        "uplink": "uplink2",
        "vdsName": "sfo-w01-cl01-vds01"
      }
    ]
  },
  "id": "<ESXi host 4 ID>",
  "licenseKey": "<license key>"
}
```

```
],  
  "isEdgeClusterConfiguredForMultiAZ": <true, if the cluster hosts an NSX Edge  
cluster; false, if the cluster does not host an NSX Edge cluster>,  
  "networkSpec": {  
    "networkProfiles": [  
      {  
        "isDefault": false,  
        "name": "sfo-w01-az2-nsx-np01",  
        "nsxtHostSwitchConfigs": [  
          {  
            "ipAddressPoolName": "sfo-w01-az2-host-ip-pool01",  
            "uplinkProfileName": "sfo-w01-az2-host-uplink-profile01",  
            "vdsName": "sfo-w01-cl01-vds01",  
            "vdsUplinkToNsxUplink": [  
              {  
                "nsxUplinkName": "uplink-1",  
                "vdsUplinkName": "uplink1"  
              },  
              {  
                "nsxUplinkName": "uplink-2",  
                "vdsUplinkName": "uplink2"  
              }  
            ]  
          }  
        ]  
      }  
    ]  
  },  
  "nsxClusterSpec": {  
    "ipAddressPoolsSpec": [  
      {  
        "description": "WLD01 AZ2 Host TEP Pool",  
        "name": "sfo-w01-az2-host-ip-pool01",  
        "subnets": [  

```

```
{
  "cidr": "172.16.44.0/24",
  "gateway": "172.16.44.253",
  "ipAddressPoolRanges": [
    {
      "end": "172.16.44.200",
      "start": "172.16.44.10"
    }
  ]
}
],
"uplinkProfiles": [
  {
    "name": "sfo-w01-az2-host-uplink-profile01",
    "teamings": [
      {
        "activeUplinks": [
          "uplink-1",
          "uplink-2"
        ],
        "name": "DEFAULT",
        "policy": "LOADBALANCE_SRCID",
        "standByUplinks": []
      }
    ],
    "transportVlan": 1644
  }
]
},
```

```

"witnessSpec": {
  "fqdn": "sfo-w01-cl01-vsw01.sfo.rainpole.io",
  "vsanCidr": "172.17.11.0/24",
  "vsanIp": "172.17.11.219"
},
"witnessTrafficSharedWithVsanTraffic": false
}
}

```

NOTE

Replace the example values in the JSON file with the correct values for your environment.

- d) Click **Execute**.
- e) On the confirmation dialog box, click **Continue**.

Configure NSX for availability zone 2.

NSX Configuration for Availability Zone 2

To provide the necessary networking services for fail-over of SDDC components from availability zone 1 to availability zone 2 in the management domain, you configure NSX for availability zone 2.

Configure IP Prefixes in the Tier-0 Gateway for Availability Zone 2

You configure default and any IP prefixes on the tier-0 gateway to permit access to route advertisement by any network and by the 0.0.0.0/0 network. These IP prefixes are used in route maps to prepend a path to one or more autonomous systems (AS-path prepend) for BGP neighbors and to configure local-reference on the learned default-route for BGP neighbors in availability zone 2.

1. In a web browser, log in to NSX Manager for the management or workload domain to be stretched at `https://nsx_manager_fqdn/login.jsp?local=true`.
2. On the main navigation bar, click **Networking**.
3. In the navigation pane, click **Tier-0 gateways**.
4. Select the gateway and from the ellipsis menu, click **Edit**.
5. Create the Any IP prefix list.
 - a) Expand the **Routing** section and in the **IP prefix list** section, click **Set**.
 - b) In the **Set IP prefix list** dialog box, click **Add IP prefix list**.
 - c) Enter *Any* as the prefix name and under **Prefixes**, click **Set**.
 - d) In the **Set prefixes** dialog box, click **Add Prefix** and configure the following settings.

Setting	Value
Network	any

Table continued on next page

Continued from previous page

Setting	Value
Action	Permit

- e) Click **Add** and then click **Apply**.
6. Repeat step 5 to create the default route IP prefix set with the following configuration.

Setting	Value
Name	Default Route
Network	0.0.0.0/0
Action	Permit

7. On the **Set IP prefix list** dialog box, click **Close**.

Configure Route Maps in the Tier-0 Gateway for Availability Zone 2

To define which routes are redistributed in the domain, you configure route maps in the tier-0 gateway.

- On the NSX Manager main navigation bar, click **Networking**.
- In the navigation pane, click **Tier-0 gateways**.
- Select the gateway, and from the ellipsis menu, click **Edit**.
- Create a route map for traffic incoming to availability zone 2.
 - Expand the **Routing** section and in the **Route maps** section, click **Set**.
 - In the **Set route maps** dialog box, click **Add route map**.
 - Enter a name for the route map.
For example, `rm-in-az2`.
 - In the **Match criteria** column, click **Set**.
 - On the **Set match criteria** dialog box, click **Add match criteria** and configure the following settings.

Setting	Value for Default Route	Value for Any
Type	IP Prefix	IP Prefix
Members	Default Route	Any
Local Preference	80	90
Action	Permit	Permit

- Click **Add** and then click **Apply**.
 - In the **Set route maps** dialog box, click **Save**.
5. Repeat step 4 to create a route map for outgoing traffic from availability zone 2 with the following configuration.

Setting	Value
Route map name	rm-out-az2
Type	IP Prefix

Table continued on next page

Continued from previous page

Setting	Value
Members	Any
As Path Prepend	<i>bgp_asn</i>
Local Preference	100
Action	Permit

6. In the **Set route maps** dialog box, click **Close**.

Configure BGP in the Tier-0 Gateway for Availability Zone 2

To enable fail-over from availability zone 1 to availability zone 2, you configure BGP neighbors on the tier-0 gateway in the management or workload domain to be stretched. You add route filters to configure `localpref` on incoming traffic and `prepend of AS` on outgoing traffic.

You configure two BGP neighbors with route filters for the uplink interfaces in availability zone 2.

Table 191: BGP Neighbors for Availability Zone 2

Setting	BGP Neighbor 1	BGP Neighbor 2
IP address	<i>ip_bgp_neighbor1</i>	<i>ip_bgp_neighbor2</i>
BFD	Deactivated	Deactivated
Remote AS	<i>asn_bgp_neighbor1</i>	<i>asn_bgp_neighbor2</i>
Hold downtime	12	12
Keep alive time	4	4
Password	<i>bgp_password</i>	<i>bgp_password</i>

Table 192: Route Filters for BGP Neighbors for Availability Zone 2

Setting	BGP Neighbor 1	BGP Neighbor 2
IP Address Family	IPV4	IPV4
Activated	Activated	Activated
Out Filter	rm-out-az2	rm-out-az2
In Filter	rm-in-az2	rm-in-az2
Maximum Routes	-	-

1. On the NSX Manager main navigation bar, click **Networking**.
2. In the navigation pane, click **Tier-0 gateways**.
3. Select the gateway and from the ellipsis menu, click **Edit**.
4. Add the uplink interfaces to the NSX Edge nodes.
 - a) Expand **BGP** and in the **BGP neighbors** section, click **2**.
 - b) In the **Set BGP neighbors** dialog box, click **Add BGP neighbor** and configure the following settings.

Setting	Value
IP address	<i>ip_bgp_neighbor1</i>
BFD	Deactivated NOTE Activate BFD only if the network supports and is configured for BFD.
Remote AS	<i>asn_bgp_neighbor1</i>
Source addresses	Select AZ2 interfaces
Hold downtime	12
Keep alive time	4
Password	<i>bgp_password</i>

- c) In the **Route filter** section, click **Set**.
- d) In the Set route filter dialog box, click **Add route filter** and configure the following settings.

Setting	Value
IP Address Family	IPV4
Enabled	Activated
Out Filter	<i>rm-out-az2</i>
In Filter	<i>rm-in-az2</i>
Maximum Routes	-

- e) Click **Add** and then click **Apply**.
5. Repeat step 4 to configure BGP neighbor *ip_bgp_neighbor2* and the corresponding route filter.
6. On the **Tier-0 gateway** page, click **Close editing**.

Expand a Stretched Cluster in VMware Cloud Foundation

You can expand a stretched cluster by adding hosts. It is recommended that you add the same number of hosts to both availability zones for symmetry and cluster balance.

- Commission the additional hosts to VMware Cloud Foundation.

See [Commission Hosts](#).

NOTE

Make sure to select the correct network pool for the availability zone to which you are adding each host.

- Get the UIDs of the hosts you commissioned.
 - In the navigation pane, click **Developer Center > API Explorer**.
 - Under APIs for managing hosts, click **GET /v1/hosts**.
 - Click **Execute**.
 - Click **Download** to download the JSON file.
 - Open the JSON file and copy the the UIDs of the hosts.
- Get the ID of the cluster you are expanding.
 - In the API Explorer, navigate to APIs for managing clusters and click **GET /v1/clusters**.

- b) Click **Execute**.
 - c) Click **Download** to download the JSON file.
 - d) Open the JSON file and copy the the cluster ID for the cluster you are expanding.
4. Get the primary and secondary availability zone names from vCenter Server.
 - a) In a web browser, log in to the vCenter Server at https://vcenter_server_fqdn/ui.
 - b) Select **Menu > Hosts and Clusters**.
 - c) In the inventory panel, expand **vCenter Server > Datacenter**.
 - d) Select **Cluster** and then click the **Configure** tab.
 - e) Under **vSAN**, select **Fault Domains**.
 - f) Note the primary and secondary availability zone names.
 5. Prepare the JSON request body.
 - a) Click **Patch /v1/clusters/id**.
 - b) Under **ClusterUpdateSpec** field, click **Cluster Update Data ClusterUpdateSpec{ ... }**.
 - c) Click **Download** to download the JSON file.
 - d) Edit the downloaded JSON file so that it contains only the expand section similar to the example below. In the **azName** field, type the primary and secondary names you had retrieved in step 4.

NOTE

The ESXi hosts that you are adding must use the same vmnic to vSphere Distributed Switch mapping as the existing hosts in the stretched cluster. For example: If existing hosts map vmnic0 and vmnic1 to vSphere Dstributed Switch 1 and vmnic2 and vmnic3 to vSphere Distributed Switch 2, then the hosts you are adding must map the same vmnics to the same vSphere Distributed Switches.

```
{
  "clusterExpansionSpec": {
    "hostSpecs": [ {
      "id": "ESXi host 1 ID",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX",
      "azName": "primary/secondary",
      "hostNetworkSpec": {
        "vmNics": [{
          "id": "vmnic0",
          "vdsName": "<vSphere Distributed Switch 1>"
        },
        {
          "id": "vmnic1",
          "vdsName": "<vSphere Distributed Switch 2>"
        }
      ]
    }
  ]
}
```

```
}, {
  "id": "ESXi host 2 ID",
  "licenseKey": "XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX",
  "azName": "primary/secondary",
  "hostNetworkSpec": {
    "vmNics": [{
      "id": "vmnic0",
      "vdsName": "<vSphere Distributed Switch 1>"
    },
    {
      "id": "vmnic1",
      "vdsName": "<vSphere Distributed Switch 2>"
    }
  ]
}
}, {
  "id": "ESXi host 3 ID",
  "licenseKey": "XXXXXX-XXXXXX-XXXXXX-XXXXXX-XXXXXX",
  "azName": "primary/secondary",
  "hostNetworkSpec": {
    "vmNics": [{
      "id": "vmnic0",
      "vdsName": "<vSphere Distributed Switch 1>"
    },
    {
      "id": "vmnic1",
      "vdsName": "<vSphere Distributed Switch 2>"
    }
  ]
}
}, {
  "id": "ESXi host 4 ID",
```

```

    "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX",
    "azName": "primary/secondary",
    "hostNetworkSpec": {
      "vmNics": [{
        "id": "vmnic0",
        "vdsName": "<vSphere Distributed Switch 1>"
      },
      {
        "id": "vmnic1",
        "vdsName": "<vSphere Distributed Switch 2>"
      }
    ]
  }
} ]
}
}

```

6. Run the expand cluster API.
 - a) For the **ClusterUpdateSpec** field, update the cluster ID (you retrieved this in step 3) and JSON file with the payload you prepared in step 5.
 - b) Click **Execute**.
 - c) Monitor the task until it is completed.
7. If required, SSH in to each newly added host and add a static route to the vSAN network of the witness host. Also add static routes in the witness if it could not reach the vSAN network of the newly added hosts.
8. Update the value of **Host failure cluster tolerates** to the number of hosts in AZ1 after cluster expansion.
 - a) Log in to the management vCenter Server.
 - b) Select **Cluster** and click the **Configure** tab.
 - c) Under **Services**, click **vSphere Availability** and then click **Edit**.
 - d) On the **Admission Control** page of the **Edit Cluster Settings** dialog box, set host failures cluster tolerates to the number of hosts in availability zone 1 and click **OK**.

Unstretch a Cluster

This procedure describes how to unstretch a vSAN cluster.

Consolidate or delete all virtual machine snapshots on the vSAN cluster before unstretching the cluster.

1. Get the ID of the cluster you are unstretching.
 - a) In the navigation pane, click **Developer Center** > **API Explorer**.
 - b) In the API Explorer, navigate to APIs for managing clusters and click `GET /v1/clusters`.
 - c) Click **Execute**.

- d) Click **Download** to download the JSON file.
 - e) Open the JSON file and copy the the cluster ID for `SDDC-Cluster1`.
2. Prepare the JSON request body.
 - a) Click **Patch /v1/clusters/id**.
 - b) Under **ClusterUpdateSpec** field, click **Cluster Update Data ClusterUpdateSpec{ ... }**.
 - c) Click **Download** to download the JSON file.
 - d) Edit the downloaded JSON file so that it contains only the unstretch information similar to the example below.

```
{ "clusterUnstretchSpec": {} }
```

3. Run the unstretch cluster API.
 - a) For the **ClusterUpdateSpec** field, update the cluster UID (you retrieved this in step 1) and unstretch JSON file with the payload you prepared in step 2.
 - b) Click **Execute**.

The unstretch cluster task is displayed in the SDDC Manager task panel.
 - c) Monitor the unstretch cluster task till it is completed.

All hosts from AZ2 are removed from the unstretched cluster and the cluster is converted to standard vSAN cluster.

Replace a Failed Host in a Stretched Cluster

If a host or host component in a stretched cluster fails, it is recommended that you replace the host with a new host.

- Image the replacement host with the same ESXi version as the other hosts in the cluster.
 - Check the health of the cluster.
1. Get the ID of the host to be removed.
 - a) In the navigation pane, click **Developer Center > API Explorer**.
 - b) Under APIs for managing hosts, click **GET /v1/hosts**.
 - c) Click **Execute**.
 - d) Click **Download** to download the JSON file.
 - e) Open the JSON file and copy the the ID of the host to be removed.
 2. Get the ID of the cluster from where the host is to be removed.
 - a) In the API Explorer, navigate to APIs for managing clusters and click `GET /v1/clusters`.
 - b) Click **Execute**.
 - c) Click **Download** to download the JSON file.
 - d) Open the JSON file and copy the the cluster ID.
 3. Prepare the JSON request body.
 - a) Click **Patch /v1/clusters/id**.
 - b) Under **ClusterUpdateSpec**, click **Cluster Update Data ClusterUpdateSpec{ ... }**.
 - c) Click **Download** to download the JSON file.
 - d) Edit the JSON file so that it contains only the compact section similar to the example below.

```
{
  "clusterCompactionSpec": {
    "hosts": [ {
      "id": "ESXi host 1 ID"
    } ]
  }
}
```

```

    }, {
      "id": "ESXi host 2 ID"
    }, {
      "id": "ESXi host 3 ID"
    } ]
  }
}

```

4. Run the compact cluster API.
 - a) In the `id` field, replace the values with the host IDs you retrieved in step 1.
 - b) Click **Execute**.
 - c) Monitor the task till it is completed.
5. Decommission the host to be removed.
See [Decommission Hosts](#).
6. Commission the replacement host to the same network pool as the removed host.
See [Commission Hosts](#).
7. Expand the cluster to add the commissioned host to the cluster. See [Expand a Stretched Cluster in](#) .
8. If required, SSH in to each newly added host and add a static route to the vSAN network of the witness host. Also add static routes in the witness if it could not reach the vSAN network of the newly added hosts.

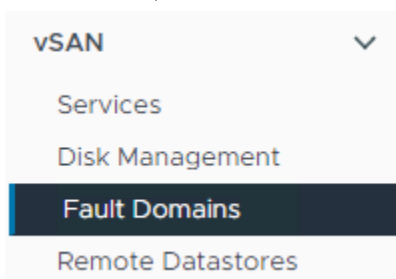
vSAN automatically rebuilds the stretch cluster.

Change the vSAN Witness Host in a Stretched Cluster

You can replace or change the vSAN witness host for a stretched cluster using the vSphere Client without affecting life cycle operations of SDDC Manager.

Verify that the vSAN witness host is not in use by another cluster, has a VMkernel configured for vSAN traffic, and has no vSAN partitions on its disks.

1. In a web browser, log in to vCenter Server at `https://vcenter_server_fqdn/ui`.
2. Select **Menu > Hosts and Clusters**.
3. In the inventory panel, expand **vCenter Server > Datacenter**.
4. Select the stretched cluster and click **Configure**.
5. Under vSAN, click **Fault Domains**.



6. Click the **Change** button.
7. Select a new host to use as the vSAN witness host and click **Next**.
8. Claim disks on the new witness host and click **Next**.
9. Review the configuration and click **Finish**.

Monitoring Capabilities in the VMware Cloud Foundation System

The VMware Cloud Foundation system provides built-in capabilities to help you perform effective operations monitoring, troubleshooting, performance management, infrastructure capacity planning, and compliance monitoring and auditing.

You use the built-in monitoring capabilities for these typical scenarios.

Scenario	Examples
Are the systems online?	A host or other component shows a failed or unhealthy status.
Why did a storage drive fail?	Hardware-centric views spanning inventory, configuration, usage, and event history to provide for diagnosis and resolution.
Is the infrastructure meeting tenant service level agreements (SLAs)?	Analysis of system and device-level metrics to identify causes and resolutions.
At what future time will the systems get overloaded?	Trend analysis of detailed system and device-level metrics, with summarized periodic reporting.
What person performed which action and when?	History of secured user actions, with periodic reporting. Workflow task history of actions performed in the system.

Viewing Tasks and Task Details

From SDDC Manager UI, you can access all tasks. By default, the Dashboard displays the Recent Tasks widget, providing general information at a glance about the most recent tasks. A task is a unit of work or a series of subtasks that perform an overall goal, such as creating a workload domain.

In addition to the most recent tasks, you can view and search for all tasks by clicking **View All Tasks** at the bottom of the Recent Tasks widget. This opens the Tasks panel.

NOTE

For more information about controlling the widgets that appear on the Dashboard page of SDDC Manager UI, see [Tour of the User Interface](#).

Viewing and Filtering Task Details

The Tasks panel provides a high level view all tasks, displaying the descriptive task name, task status (for example, running, succeeded, or failed), and the timestamp for the last change in task status. You can also filter and search the task information as follows:

- Search tasks by clicking the filter icon in the Task column header and entering a search string.
- Filter tasks by status by clicking the filter icon in Status column. Select by category **All**, **Failed**, **Successful**, **Running**, or **Pending**.

NOTE

Each category also displays the number of tasks with that status.

- Clear all filters by clicking **Reset Filter** at the top of the Tasks panel.
- Click **Refresh** to refresh the task list.

NOTE

You can also sort the table by the contents of the Status and Last Occurrence columns.

Managing Tasks and Subtask Details

Expand a task to view details including the subtasks that comprise the task and their individual statuses.

- If a task is in a Failed state, you can also attempt to restart it by clicking **Restart Task**.

NOTE

Not all tasks are restartable.

- If a task is in a Failed state, click on the icon next to the Failed status to view a detailed report on the cause.
- To view subtasks and their details, click **View Subtasks**.

NOTE

You can filter subtasks in the same way you filter tasks.

NOTE

You can also sort the table by the contents of the Status and Last Occurrence columns.

Resizing the Task Panel

Use the icons on the task panel to increase or decrease the panel size, or to close or reopen it.

API Activity Logging

When you invoke APIs or log in to or log out from the SDDC Manager UI, VMware Cloud Foundation creates activity log files that track the request. Activity logs can be used to analyze the pattern of user actions and gather metrics.

The following logs are available on the SDDC Manager appliance:

Log Name	Location
sddc-manager-ui-activity.log	/var/log/vmware/vcf/sddc-manager-ui-app
domainmanager-activity.log	/var/log/vmware/vcf/domainmanager
operationsmanager-activity.log	/var/log/vmware/vcf/operationsmanager
lcm-activity.log	/var/log/vmware/vcf/lcm
vcf-commonsvcs-activity.log	/var/log/vmware/vcf/commonsvcs

Activity Log Structure

All activity logs use the following JSON schema:

```
{
  "timestamp": "", "username": "", "clientIP": "", "userAgent": "", "api": "", "httpMethod": "",
  "httpStatus": "", "operation": "", "remoteIP": ""
}
```

Activity Log Example

The following example is from the domainmanager-activity.log:

```
{"username": "administrator@vsphere.local", "timestamp": "2022-01-19T16:59:01.919Z ",
"client IP": "10.0.0.253", "userAgent": "Mozilla/5.0 (Windows NT 6.3; Win 64; x64)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/97.0.4692.71 Safari/537.36", "api": "/"
```

```
domainmanager/vl/vra/domains", "httpMethod": "GET", "httpStatus": 200, "operation": "Gets VMware Aria Automation integration status for workload domains", "remote IP": "127.0.0.1"}
```

- **username:** The username of the system from which the API request is triggered. For example: "administrator@vsphere.local".
- **timestamp:** Date and time of the operation performed in the UTC format "YYYY-MM-DD'THH:MM:SS.SSSXXX". For example: "2022-01-19T16:59:01.9192".
- **client IP:** The IP address of the user's system. For example: "10.0.0.253".
- **userAgent:** The user's system information such as the web browser name, web browser version, operating system name, and operating system architecture type. For example: "Mozilla/5.0 (Windows NT 6.3; Win 64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/97.0.4692.71 Safari/537.36".
- **api:** The API invoked to perform the operation. For example: "/domainmanager/vl/vra/domains".
- **httpMethod:** HTTP method of the REST API. For example: "GET".
- **httpStatus:** The response code received after invoking the API. For example: 200.
- **operation:** The operation or activity that was performed. For example: "Gets VMware Aria Automation integration status for workload domains".
- **remoteIP:** remoteIP of the request initiator. For example: "127.0.0.1"

Activity Logs Retention Policy

Log files are rolled over daily to a file using the following naming format: <service-name>.<YYYY>-<MM>-<DD>.0.log.gz. For example: domainmanager.2022-01-22.0.log.gz.

The log history is stored for 30 days. The maximum file size of the log retention file is set to 100 MB.

Log Analysis

You can perform log aggregation and analysis by integrating VMware Aria Operations for Logs with VMware Cloud Foundation. For more information, see [Implementation of Intelligent Logging and Analytics for VMware Cloud Foundation](#).

Updating VMware Cloud Foundation DNS and NTP Servers

If you need to update the DNS or NTP servers that VMware Cloud Foundation uses, you can update the servers using the SDDC Manager UI.

When you initially deploy VMware Cloud Foundation, you complete the deployment parameter workbook to provide the system with the information required for bring-up. This includes up to two DNS servers and up to two NTP servers. You can reconfigure these settings at a later date, using the SDDC Manager UI.

Update DNS Server Configuration

Use this procedure to update the DNS server configuration across VMware Cloud Foundation components.

- Verify that both forward and reverse DNS resolution are functional for each VMware Cloud Foundation component using the updated DNS server information.
- Verify that the new DNS server is reachable from each of the VMware Cloud Foundation components.
- Verify all VMware Cloud Foundation components are reachable from SDDC Manager.
- Verify that all VMware Cloud Foundation components are in an `Active` state.

SDDC Manager uses DNS servers to provide name resolution for the components in the system. When you update the DNS server configuration, SDDC Manager performs DNS configuration updates for the following components:

- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX Managers
- NSX Edge nodes
- VMware Aria Suite Lifecycle
- VMware Aria Operations for Logs

If the update fails, SDDC Manager rolls back the DNS settings for the failed component. Fix the underlying issue and retry the update starting with the failed component.

NOTE

There is no rollback for VMware Aria Suite Lifecycle. Check the logs, resolve any issues, and retry the update.

Updating the DNS server configuration can take some time to complete, depending on the size of your environment. Schedule DNS updates at a time that minimizes the impact to the system users.

This procedure uses the SDDC Manager UI.

1. In the SDDC Manager UI, click **Administration > Network Settings**.
2. On the **Network Settings** page, click the **DNS Configuration** tab.
3. To update the DNS servers, click **Edit**.
4. Update the DNS configuration.
 - a) Expand the **Overview** section, and click **Next**.
 - b) Expand the **Prerequisites** section, and click **Next**.
 - c) Expand the **Edit DNS configuration** section, update the **Primary DNS server** and **Alternative DNS server**, and click **Save**.

NOTE

Alternative DNS server is optional.

Update NTP Server Configuration

Use this procedure to update the NTP server configuration across VMware Cloud Foundation components.

- Verify the new NTP server is reachable from the VMware Cloud Foundation components.
- Verify the time skew between the new NTP servers and the VMware Cloud Foundation components is less than 5 minutes.
- Verify all VMware Cloud Foundation components are reachable from SDDC Manager.
- Verify all VMware Cloud Foundation components are in an *Active* state.

SDDC Manager uses NTP servers to synchronize time between the components in the system. You must have at least one NTP server. When you update the NTP server configuration, SDDC Manager performs NTP configuration updates for the following components:

- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX Managers
- NSX Edge nodes
- VMware Aria Suite Lifecycle
- VMware Aria Operations for Logs
- VMware Aria Operations
- VMware Aria Automation

If the update fails, SDDC Manager rolls back the NTP settings for the failed component. Fix the underlying issue and retry the update starting with the failed component.

NOTE

There is no rollback for the VMware Aria Suite Lifecycle. Check the logs, resolve any issues, and retry the update.

Updating the NTP server configuration can take some time to complete, depending on the size of your environment. Schedule NTP updates at a time that minimizes the impact to the system users. This procedure uses the SDDC Manager UI.

1. In the SDDC Manager UI, click **Administration > Network Settings**.
2. On the **Network Settings** page, click the **NTP Configuration** tab.
3. To update the NTP servers, click **Edit**.
4. Update the NTP configuration.
 - a) Expand the **Overview** section, and click **Next**.
 - b) Expand the **Prerequisites** section, and click **Next**.
 - c) Expand the **Edit NTP configuration** section, update the **NTP server**, and click **Save**.

Supportability and Serviceability (SoS) Utility

The SoS utility is a command-line tool that you can use to run health checks, collect logs for VMware Cloud Foundation components, and so on.

To run the SoS utility, SSH in to the SDDC Manager appliance using the `vcf` user account. For basic operations, enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-1 --option-2 --option-3 ... --option-n
```

To list the available command options, use the `--help` long option or the `-h` short option.

```
sudo /opt/vmware/sddc-support/sos --help
```

```
sudo /opt/vmware/sddc-support/sos -h
```

NOTE

You can specify options in the conventional GNU/POSIX syntax, using `--` for the long option and `-` for the short option.

For privileged operations, enter `su` to switch to the root user, and navigate to the `/opt/vmware/sddc-support` directory and type `./sos` followed by the options required for your desired operation.

SoS Utility Options

This section lists the specific options you can use with the SoS utility.

For information about collecting log files using the SoS utility, see [Collect Logs for Your System](#).

SoS Utility Help Options

Use these options to see information about the SoS utility itself. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

Option	Description
--help -h	Provides a summary of the available SoS utility options
--version -v	Provides the SoS utility's version number.

SoS Utility Generic Options

These are generic options for the SoS utility. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

Option	Description
--history	Displays the last 20 SoS operations performed.
--force	Allows SoS operations to be performed while workflows are running. NOTE It is recommended that you do not use this option.
--configure-sftp	Configures SFTP for logs.
--setup-json <i>SETUPJSON</i>	Custom setup-json file for log collection. SoS prepares the inventory automatically based on the environment where it is running. If you want to collect logs for a pre-defined set of components, you can create a <code>setup.json</code> file and pass the file as input to SoS. A sample JSON file is available on the SDDC Manager appliance at <code>/opt/vmware/sddc-support/setup.sample.json</code> .
--log-folder <i>LOGFOLDER</i>	Specifies the name of the log directory.
--log-dir <i>LOGDIR</i>	Specifies the directory to store the logs.
--enable-stats	Activate SoS execution stats collection.
--debug-mode	Runs the SoS utility in debug mode.
--zip	Creates a zipped TAR file for the output.
--short	Display detailed health results only for failures and warnings.
--domain-name <i>DOMAINNAME</i>	Specify the name of the workload domain name on which to perform the SoS operation. To run the operation on all workload domains, specify <code>--domain-name ALL</code> . NOTE If you omit the <code>--domain-name</code> flag and workload domain name, the SoS operation is performed only on the management domain. You can combine <code>--domain-name</code> with <code>--clusternames</code> to further limit the scope of an operation. This can be useful in a scaled environment with a large number of ESXi hosts.

Table continued on next page

Continued from previous page

Option	Description
<code>--clusternames</code> <i>CLUSTER NAMES</i>	Specify the vSphere cluster names associated with a workload domain for which you want to collect ESXi and Workload Management (WCP) logs. Enter a comma-separated list of vSphere clusters. For example, <code>--clusternames cluster1, cluster2</code> . NOTE If you specify <code>--domain-name ALL</code> then the <code>--clusternames</code> option is ignored.
<code>--skip-known-host-check</code>	Skips the specified check for SSL thumbprint for host in the known host.
<code>--include-free-hosts</code>	Collect logs for free ESXi hosts, in addition to in-use ESXi hosts.
<code>--include-precheck-report</code>	This option runs LCM upgrade prechecks and includes the LCM upgrade prechecks run report in SoS health check operations.

SoS Utility VMware Cloud Foundation Summary Options

These options provide summary details of the SDDC Manager instance, including components, services, and tasks.. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

Option	Description
<code>--get-vcf-summary</code>	Returns information about your VMware Cloud Foundation system, including CEIP, workload domains, vSphere clusters, ESXi hosts, licensing, network pools, SDDC Manager, and VCF services.
<code>--get-vcf-tasks-summary</code>	Returns information about VMware Cloud Foundation tasks, including the time the task was created and the status of the task.
<code>--get-vcf-services-summary</code>	Returns information about SDDC Manager uptime and when VMware Cloud Foundation services (for example, LCM) started and stopped.

SoS Utility Fix-It-Up Options

Use these options to manage ESXi hosts and vCenter Servers, including enabling SSH and locking down hosts. For these options, SSH in to the SDDC Manager VM using the `vcf` administrative user account, enter `su` to switch to the root user, navigate to the `/opt/vmware/sddc-support` directory, and type the following command:

```
./sos --option-name
```

NOTE

For Fix-It-Up options, if you do not specify a workload domain, the command affects only the management domain.

Option	Description
<code>--enable-ssh-esxi</code>	Applies SSH on all ESXi nodes in the specified workload domains.

Table continued on next page

Continued from previous page

Option	Description
	<ul style="list-style-type: none"> To enable SSH on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To enable SSH on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--disable-ssh-esxi</code>	<p>Deactivates SSH on all ESXi nodes in the specified workload domains.</p> <ul style="list-style-type: none"> To deactivate SSH on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To deactivate SSH on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--enable-ssh-vc</code>	<p>Applies SSH on vCenter Server in the specified workload domains.</p> <ul style="list-style-type: none"> To enable SSH on vCenter in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To enable SSH on vCenter Servers in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--disable-ssh-vc</code>	<p>Deactivates SSH on vCenter Servers in the specified workload domains.</p> <ul style="list-style-type: none"> To deactivate SSH on vCenter Server in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To deactivate SSH on vCenter Servers in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--enable-lockdown-esxi</code>	<p>Applies normal lockdown mode on all ESXi nodes in the specified workload domains.</p> <ul style="list-style-type: none"> To enable lockdown on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To enable lockdown on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--disable-lockdown-esxi</code>	<p>Deactivates normal lockdown mode on ESXi nodes in the specified workload domains.</p> <ul style="list-style-type: none"> To deactivate lockdown on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To deactivate lockdown on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--ondemand-service ONDEMANDSERVICE</code>	<p>Execute commands on ESXi hosts, vCenter Servers, or SDDC Manager entities for a given workload domain. Specify the workload domain using <code>--domain-name DOMAINNAME</code>.</p> <p>Replace <code>ONDEMANDSERVICE</code> with the path to a <code>.yaml</code> input file. (Sample file available at: <code>/opt/vmware/sddc-support/ondemand_command_sample.yaml</code>).</p> <p style="text-align: center;">WARNING Contact Broadcom Support before using this option.</p>
<code>--ondemand-service JSON file path</code>	<p>Include this flag to execute commands in the JSON format on all ESXi hosts in a workload domain. For example, <code>/opt/vmware/sddc-support/<JSON file name></code></p>

Table continued on next page

Continued from previous page

Option	Description
<code>--refresh-ssh-keys</code>	Refreshes the SSH keys.

SoS Utility Health Check Options

These SoS commands are used for checking the health status of various components or services, including connectivity, compute, storage, database, workload domains, and networks. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

A green status indicates that the health is normal, yellow provides a warning that attention might be required, and red (critical) indicates that the component needs immediate attention.

Option	Description
<code>--health-check</code>	Performs all available health checks. Can be combined with <code>--run-vsan-checks</code> . For example: <pre>sudo /opt/vmware/sddc-support/sos --health-check --run-vsan-checks</pre>
<code>--connectivity-health</code>	Performs connectivity checks and validations for SDDC resources (NSX Managers, ESXi hosts, vCenter Servers, and so on). This check performs a ping status check, SSH connectivity status check, and API connectivity check for SDDC resources.
<code>--services-health</code>	Performs a services health check to confirm whether services within the SDDC Manager (like Lifecycle Management Server) and vCenter Server are running.
<code>--compute-health</code>	Performs a compute health check, including ESXi host licenses, disk storage, disk partitions, and health status.
<code>--storage-health</code>	Performs a check on the vSAN disk health of the ESXi hosts and vSphere clusters. Can be combined with <code>--run-vsan-checks</code> . For example: <pre>sudo /opt/vmware/sddc-support/sos --storage-health --run-vsan-checks</pre>
<code>--run-vsan-checks</code>	This option cannot be run on its own and must be combined with <code>--health-check</code> or <code>--storage-health</code> . Runs a VM creation test to verify the vSAN cluster health. Running the test creates a virtual machine on each host in the vSAN cluster. The test creates a VM and deletes it. If the VM creation and deletion tasks are successful, assume that the vSAN cluster components are working as expected and the cluster is functional. NOTE You must not conduct the proactive test in a production environment as it creates network traffic and impacts the vSAN workload.
<code>--ntp-health</code>	Verifies whether the time on the components is synchronized with the NTP server in the SDDC Manager appliance. It also ensures that the hardware and software time stamp of ESXi hosts are within 5 minutes of the SDDC Manager appliance.
<code>--dns-health</code>	Performs a forward and reverse DNS health check.

Table continued on next page

Continued from previous page

Option	Description
<code>--general-health</code>	Checks ESXi for error dumps and gets NSX Manager and cluster status.
<code>--certificate-health</code>	Verifies that the component certificates are valid and when they are expiring. <ul style="list-style-type: none"> • GREEN: Certificate expires in more than 30 days. • YELLOW: Certificate expires in 15-30 days. • RED: Certificate expires in less than 15 days.
<code>--get-host-ips</code>	Returns host names and IP addresses of ESXi hosts.
<code>--get-inventory-info</code>	Returns inventory details for the VMware Cloud Foundation components, such as vCenter ServerNSX, SDDC Manager, and ESXi hosts. Optionally, add the flag <code>--domain-name ALL</code> to return details for all workload domains.
<code>--password-health</code>	Checks the status of passwords across VMware Cloud Foundation components. It lists components with passwords managed by VCF, the date a password was last changed, the password expiration date, and the number of days until expiration. <ul style="list-style-type: none"> • GREEN: Password expires in more than 15 days. • YELLOW: Password expires in 5-15 days. • RED: Password expires in less than 5 days.
<code>--hardware-compatibility-report</code>	Validates ESXi hosts and vSAN devices and exports the compatibility report.
<code>--version-health</code>	This operation checks the version of BOM components (vCenter Server, NSX, ESXi, and SDDC Manager). It compares the SDDC Manager inventory, the actual installed BOM component version, and the BOM component versions to detect any drift.
<code>--json-output-dir JSONDIR</code>	Outputs the results of any health check as a JSON file to the specified directory, JSONDIR.

Example Health Check Commands:

- Check the password health on the management domain only:
`./sos --password-health`
- Check the connectivity health for all workload domains:
`./sos --connectivity-health --domain-name ALL`
- Check the DNS health for the workload domain named `sfo-w01`:
`./sos --dns-health --domain-name sfo-w01`

Collect Logs for Your VMware Cloud Foundation System

Use the SoS utility to collect the logs for various software components in the system.

Use these options when retrieving support logs from your environment's various components.

- If you run the SoS utility from SDDC Manager without specifying any component-specific options, the SoS tool collects SDDC Manager, API, and VMware Cloud Foundation summary logs. To collect all logs, use the `--collect-all-logs` options.

NOTE

SoS log collection may time out after 60 minutes, which could be an issue with large workload domains. If the SoS utility does time out, collect component-specific logs or limit log collection to specific clusters using the options described below.

- If you run the SoS utility from Cloud Builder without specifying any component-specific options, the SoS tool collects SDDC Manager, API, and Cloud Builder logs.
- To collect logs for a specific component, run the utility with the appropriate options. For example, the `--domain-name` option is important. If omitted, the SoS operation is performed only on the management domain. See [SoS Utility Options](#).

After running the SoS utility, you can examine the resulting logs to troubleshoot issues, or provide to VMware Technical Support if requested. VMware Technical Support might request these logs to help resolve technical issues when you have submitted a support request. The diagnostic information collected using the SoS utility includes logs for the various VMware software components and software products deployed in your VMware Cloud Foundation environment.

Table 193: SoS Utility Log File Options

Option	Description
<code>--esx-logs</code>	Collects logs from the ESXi hosts only. Logs are collected from each ESXi host available in the deployment.
<code>--vc-logs</code>	Collects logs from the vCenter Server instances only. Logs are collected from each vCenter server available in the deployment.
<code>--sddc-manager-logs</code>	Collects logs from the SDDC Manager only. <code>sddc<timestamp>.tgz</code> contains logs from the SDDC Manager file system's <code>etc</code> , <code>tmp</code> , <code>usr</code> , and <code>var</code> partitions.
<code>--vxrail-manager-logs</code>	Collects logs from VxRail Manager instances only.
<code>--psc-logs</code>	Collects logs from the Platform Services Controller instances only.
<code>--nsx-logs</code>	Collects logs from the NSX Manager and NSX Edge instances only.
<code>--wcp-logs</code>	Collects logs from Workload Management clusters only.
<code>--vrealize-logs</code>	Collects logs from VMware Aria Suite Lifecycle.
<code>--no-clean-old-logs</code>	Use this option to prevent the utility from removing any output from a previous collection run. By default, before writing the output to the directory, the utility deletes the prior run's output files that might be present. If you want to retain the older output files, specify this option.
<code>--test</code>	Collects test logs by verifying the files.
<code>--no-health-check</code>	Skips the health check executed as part of log collection.
<code>--api-logs</code>	Collects output from REST endpoints for SDDC Manager inventory and LCM.
<code>--rvc-logs</code>	Collects logs from the Ruby vSphere Console (RVC) only. RVC is an interface for ESXi and vCenter. NOTE If the Bash shell is not enabled in vCenter Server, RVC log collection will be skipped . NOTE RVC logs are not collected by default with <code>./sos</code> log collection. You must enable RVC to collect RVC logs.
<code>--vm-screenshots</code>	Collects all VM screenshots.
<code>--system-debug-logs</code>	Collects system logs to help with debugging uncommon issues.

Table continued on next page

Continued from previous page

Option	Description
<code>--collect-all-logs</code>	Collects logs for all components, except Workload Management and system debug logs. By default, logs are collected for the management domain components. To collect logs for all workload domain, specify <code>--domain-name ALL</code> . To collect logs for a specific workload domain, specify <code>--domain-name domain_name</code> .
<code>--log-dir LOGDIR</code>	Specifies the directory to store the logs.
<code>--log-folder LOGFOLDER</code>	Specifies the name of the log directory.
<code>--domain-name DOMAINNAME</code>	Specify the name of the workload domain name on which the SoS operation is to be performed. To run the operation on all domains, specify <code>--domain-name ALL</code> . NOTE If you omit the <code>--domain-name</code> flag and domain name, the SoS operation is performed only on the management domain.

1. Using SSH, log in to the SDDC Manager appliance as the `vcf` user.
2. To collect the logs, run the SoS utility without specifying any component-specific options.

```
sudo /opt/vmware/sddc-support/sos
```

Enter the `vcf` password when prompted.

To collect logs for a specific component, run the utility with the appropriate options.

```
sudo /opt/vmware/sddc-support/sos --option-name
```

NOTE

By default, before writing the output to the directory, the utility deletes the prior run's output files that might be present. If you want to retain the older output files, specify the `--no-clean-old-logs` option.

If you do not specify the `--log-dir` option, the utility writes the output to the `/var/log/vmware/vcf/sddc-support` directory in the SDDC Manager appliance

The utility collects the log files from the various software components in all of the racks and writes the output to the directory named in the `--log-dir` option. Inside that directory, the utility generates output in a specific directory structure.

```
vcf@sddc-manager [ ~ ]$ sudo /opt/vmware/sddc-support/sos --domain-name MGMT --skip-known-host-check --log-dir /tmp/new
```

```
[sudo] password for vcf
```

```
Welcome to Supportability and Serviceability(SoS) utility!
```

```
Performing SoS operation for MGMT domain components
```

```
Logs : /tmp/new/sos-2019-09-03-21-04-40-11793
```

Log file : /tmp/new/sos-2019-09-03-21-04-40-11793/sos.log

Log Collection completed successfully for : [HEALTH-CHECK, SDDC-MANAGER, NSX_MANAGER, API-LOGS, ESX, VMS_SCREENSHOT, VCENTER-SERVER, VCF-SUMMARY]

Change to the output directory to examine the collected log files.

Component Log Files Collected by the SoS Utility

The SoS utility writes the component log files into an output directory structure within the file system of the SDDC Manager instance in which the command is initiated, for example:

```
vcf@sddc-manager [ ~ ]$ sudo /opt/vmware/sddc-support/sos
[sudo] password for vcf
Welcome to Supportability and Serviceability(SoS) utility!
Performing SoS operation for MGMT domain components
Logs : /var/log/vmware/vcf/sddc-support/sos-2019-09-03-20-55-41-10053
Log file : /var/log/vmware/vcf/sddc-support/sos-2019-09-03-20-55-41-10053/sos.log
NOTE : The Health check operation was invoked without --skip-known-host-check, and so will skip Connectivity Health, Password Health and Certificate Health Checks because of security reasons.
```

```
Log Collection completed successfully for : [HEALTH-CHECK, SDDC-MANAGER, NSX_MANAGER, API-LOGS, ESX, VMS_SCREENSHOT, VCENTER-SERVER, VCF-SUMMARY]
```

esx Directory Contents

In each rack-specific directory, the `esx` directory contains the following diagnostic files collected for each ESXi host in the rack:

File	Description
<code>esx-FQDN.tgz</code>	Diagnostic information from running the <code>vm-support</code> command on the ESXi host. An example file is <code>esx-esxi-1.vrack.vsphere.local.tgz</code> .
<code>SmartInfo-FQDN.txt</code>	S.M.A.R.T. status of the ESXi host's hard drive (Self-Monitoring, Analysis, and Reporting Technology). An example file is <code>SmartInfo-esxi-1.vrack.vsphere.local.txt</code> .
<code>vsan-health-FQDN.txt</code>	VMware vSAN cluster health information from running the standard command <code>python /usr/lib/vmware/vsan/bin/vsan-health-status.py</code> on the ESXi host. An example file is <code>vsan-health-esxi-1.vrack.vsphere.local.txt</code> .

nsx Directory Contents

In each rack-specific directory, the `nsx` directory contains the diagnostic information files collected for the NSX Managers and NSX Edge instances deployed in that rack.

The number of files in this directory depends on the number of NSX Manager and NSX Edge instances that are deployed in the rack. In a given rack, each management domain has a cluster of three NSX Managers. The first VI workload domain has an additional cluster of three NSX Managers. Subsequent VI workload domains can deploy their own NSX Manager cluster, or use the same cluster as an existing VI workload domain. NSX Edge instances are optional.

File	Description
VMware-NSX-Manager-tech-support- <i>nsxmanagerIPAddr</i> .tar.gz	Standard NSX Manager compressed support bundle, generated using the NSX API POST <code>https://<i>nsxmanagerIPAddr</i>/api/1.0/appliance-management/techsupportlogs/NSX</code> , where <i>nsxmanagerIPAddr</i> is the IP address of the NSX Manager instance. An example is VMware-NSX-Manager-tech-support-10.0.0.8.tar.gz.
VMware-NSX-Edge-tech-support- <i>nsxmanagerIPAddr-edgeId</i> .tgz NOTE This information is only collected if NSX Edges are deployed.	Standard NSX Edge support bundle, generated using the NSX API to query the NSX Edge support logs: GET <code>https://<i>nsxmanagerIPAddr</i>/api/4.0/edges/<i>edgeId</i>/techsupportlogs</code> , where <i>nsxmanagerIPAddr</i> is the IP address of the NSX Manager instance and <i>edgeID</i> identifies the NSX Edge instance. An example is VMware-NSX-Edge-tech-support-10.0.0.7-edge-1.log.gz.

vc Directory Contents

In each rack-specific directory, the `vc` directory contains the diagnostic information files collected for the vCenter Server instances deployed in that rack.

The number of files in this directory depends on the number of vCenter Server instances that are deployed in the rack. In a given rack, each management domain has one vCenter Server instance, and any VI workload domains in the rack each have one vCenter Server instance.

File	Description
VC- <i>vcsaFQDN</i> -vm-support.tgz	Standard vCenter Server support bundle downloaded from the vCenter Server Appliance instance having a fully qualified domain name <i>vcsaFQDN</i> . The support bundle is obtained from the instance using the standard <code>vc-support.sh</code> command.

Replacing Host Components in VMware Cloud Foundation

This section provides procedures for repairing or replacing hosts in VMware Cloud Foundation. These procedures are provided for scenarios where there is no risk of data loss, such as repairing an unassigned host or replacing a host in a non-vSAN-based VI workload domain.

These procedures are not intended for scenarios where there is a risk of data loss or involving a catastrophic failure. If there is a risk of data loss or you have experienced a catastrophic failure, before taking any steps to remediate the

situation, contact Broadcom Support to review your recovery plan. This strategy ensures that additional damage is not done while troubleshooting.

NOTE

Before performing any maintenance, review [Avoiding Unintentional Downtime](#).

For covered failure scenarios, the replacement procedure depends on the component being replaced and the condition of the component.

Avoiding Unintentional Downtime

Many outages are caused by human error. Before performing maintenance on hosts or the network and storage infrastructure on which they depend, take precautions to avoid unintentional downtime.

There are a few steps you can take to avoid unintentional downtime.

- For operations that could impact access to storage, check that you have current backups for the VMs in the vSphere cluster, and if not, take a backup before proceeding. These operations include maintenance on hosts in a vSAN cluster, datastore mount points for external storage, and storage-array LUN masks.
- For operations that may take compute capacity offline, check that there is sufficient capacity available to continue running the VMs if the host being repaired cannot be brought back online as planned.

In addition, if the maintenance involves a host in a vSphere cluster, before proceeding with the maintenance, check the vCenter Server and NSX Managers associated with the host for any alerts that indicate a problem beyond the one you are planning to fix. If there are any alerts, address them first.

Replacing Components of a Host Running in Degraded Mode

The procedures for replacing components of hosts in a degraded state depend on whether the host is assigned or unassigned. An assigned host belongs to a workload domain. An unassigned host has been commissioned, but is not assigned to a workload domain.

Before proceeding, review the guidance provided in the [Avoiding Unintentional Downtime](#).

These procedures apply to the following components:

- CPU
- Memory
- Out-of-Band Management
- Power supply

Replace Components of an Assigned Host Running in Degraded Mode

This procedure shows you how to replace the components of an assigned host running in degraded mode.

- Verify that the host is operational and is accessible using the VMware Host Client.
- Verify that the Management, vSAN, and vMotion networks are available on the host. This can be viewed through the **Inventory > Hosts** page.
- Verify that the HDD and SSD disks on the host are in a good state.
- Verify that there are no alerts reported in vCenter Server for the host's cluster, and, if the cluster is a vSAN cluster, verify there are no vSAN health alerts.

1. In the navigation pane, click **Inventory > Workload Domains**.
2. In the Domain column, click the workload domain name where the host is assigned.
3. Click the **Services** tab.
4. Click the vCenter Server launch link.
5. Place the host in maintenance mode using one of the following methods, depending on the type of principal storage.

Principal Storage	Steps
vSAN	<ol style="list-style-type: none"> 1. Select the vSphere cluster the host is part of, then click Monitor tab. 2. Under vSAN, click Data Migration Pre-check. 3. Using the Pre-check data migration of drop-down, select the host. 4. Using the vSAN data migration drop-down, select Full data migration and click Pre-Check. 5. If the test returns as a success, click Enter Maintenance Mode. 6. On the Enter maintenance mode dialog, click OK.
Non-vSAN	<ol style="list-style-type: none"> 1. Expand the vSphere cluster the host is part of, right-click the affected host and click Maintenance Mode > Enter Maintenance Mode. 2. On the Enter maintenance mode dialog, click OK.

6. Right-click the affected host and select **Shutdown**.

7. Pull the host out of the physical rack.

Note the ports on the switches it was connected to.

8. Service the appropriate part following the OEM vendor documentation.

9. Put the host back in the physical rack and connect it back to the appropriate switches.

10. Power on the host.

11. In vSphere Client, right-click the host and click **Maintenance Mode > Exit Maintenance Mode**.

Replace Components of an Unassigned Host Running in Degraded Mode

This procedure shows you how to replace the components of an unassigned host that is running in degraded mode.

- Verify that the host is operational and is accessible by VMware Host Client.
- Verify that the HDD and SSD disks on the host are in a good state.

1. In the navigation pane, click **Inventory > Hosts**.

2. In the Hosts table, select the host on which you want to perform maintenance.

3. From the Action menu, click **Open in VMware Host Client** and log in.

4. Right-click the host and select **Shutdown**.

5. Pull the host out of the physical rack.

Note the ports on the switches it was connected to.

6. Service the appropriate part following the OEM vendor documentation.

7. Put the host back in the physical rack and connect it back to the appropriate switches.

8. Power on the host.

9. In SDDC Manager UI, verify that the host is available in the free pool.

Replace a Dead Host

If you need to replace a dead host, you must remove the host from the physical rack. You can then add a new host or replace the failed component on the host and add it back.

If the host is assigned to a workload domain, verify that there are at least four hosts in the management domain or at least three hosts in a VI workload domain to which the faulty host belongs. If there are fewer than the required number of hosts, contact Broadcom Support for assistance. Before proceeding, review the guidance provided in [Avoiding Unintentional Downtime](#).

This procedure applies chiefly to the following components:

- Storage controllers
- Motherboards
- Boot disks

1. If the host is in assigned to a workload domain, it must be forcibly removed.

See [Remove a Host from a vSphere Cluster in a Workload Domain](#).

2. Decommission the host.

See [Decommission Hosts](#).

3. Power off the host and remove it from the physical rack.

4. Replace and reconfigure, as follows.

- a) Replace the failed component on the host.
- b) Perform a fresh reinstall of ESXi.
- c) Commission the host.

See [Commission Hosts](#).

Replace Boot Disk on a Host

This section describes the replacement procedure for a failed boot disk on a host.

If the host is operational, verify that there are at least four hosts in the management domain or at least three hosts in a VI workload domain to which the faulty host belongs. If there are fewer than the required number of hosts, add a host to the workload domain from the free pool, if possible. If the host is not operational, see [Replace a Dead Host](#).

1. If there are dual boot disks in the host setup as RAID 1 and only one of them fails:

- See [Replacing Components of a Host Running in Degraded Mode](#) to replace the failed disk.

The RAID 1 feature will rebuild the disks as needed. For more details, refer to the OEM vendor documentation.

2. If there is a single boot disk in the host and it fails, see [Replace a Dead Host](#).

Managing Users and Groups in VMware Cloud Foundation

You can add users and groups to VMware Cloud Foundation to provide users with access to the SDDC Manager UI as well as the vCenter Server and NSX Manager instances that are deployed in your VMware Cloud Foundation system. Users can log in and perform tasks based on their assigned role.

Before you can add users and groups to VMware Cloud Foundation, you must configure an identity provider that has access to user and group data. VMware Cloud Foundation supports the following identity providers:

- vCenter Single Sign-On is vCenter Server's built-in identity provider. By default, it uses the system domain (for example, `vsphere.local`) as its identity source. You can add Active Directory over LDAP and OpenLDAP as identity sources for vCenter Single Sign-On.
- You can also use any of the following external identity providers instead of vCenter Single Sign-On:
 - Microsoft ADFS
 - Okta
 - Microsoft Entra ID (formerly known as Azure Active Directory)

Once you have configured an identity provider, you can add users and groups, and assign roles to determine what tasks they can perform from the SDDC Manager UI and VMware Cloud Foundation API.

NOTE

SDDC Manager only manages users and groups for the management SSO domain. If you created isolated VI workload domains that use different SSO domains, you must use the vSphere Client to manage users and groups for those SSO domains. Use the vSphere Client to connect to the VI workload domain's vCenter Server and then click **Administration > Single Sign On**.

In addition to user accounts, VMware Cloud Foundation includes the following accounts:

- Automation accounts for accessing VMware Cloud Foundation APIs. You can use these accounts in automation scripts.
- Local account for accessing VMware Cloud Foundation APIs when vCenter Server is down.
- Service accounts are automatically created by VMware Cloud Foundation for inter-product interaction. These are for system use only.

Configuring the Identity Provider for VMware Cloud Foundation

You can use vCenter Single Sign-On, Microsoft ADFS, Okta, or Microsoft Entra ID as the identity provider for VMware Cloud Foundation.

By default, VMware Cloud Foundation uses vCenter Single Sign-On as its identity provider and the system domain (for example, `vsphere.local`) as its identity source. You can add Active Directory over LDAP and OpenLDAP as identity sources for vCenter Single Sign-On. See [Add Active Directory over LDAP or OpenLDAP as an Identity Source for VMware Cloud Foundation](#).

You can also configure VMware Cloud Foundation to use Microsoft ADFS, Okta, or Microsoft Entra ID as an external identity provider, instead of using vCenter Single Sign-On:

- [Configure Microsoft ADFS as the Identity Provider in the SDDC Manager UI](#)
- [Configure Okta as the Identity Provider in the SDDC Manager UI](#)
- [Configure Identity Federation in VMware Cloud Foundation Using Microsoft Entra ID](#)

Add Active Directory over LDAP or OpenLDAP as an Identity Source for VMware Cloud Foundation

Users can log in to the SDDC Manager UI only if they are in a domain that has been added as a vCenter Single Sign-On identity source. vCenter Single Sign-On administrator users can add identity sources, or change the settings for identity sources that they added.

You can use identity sources to attach one or more domains to vCenter Single Sign-On. A domain is a repository for users and groups that the vCenter Single Sign-On server can use for user authentication with VMware Cloud Foundation. By default, vCenter Single Sign-On includes the system domain (for example, `vsphere.local`) as an identity source. You can add Active Directory over LDAP or an OpenLDAP directory service as identity sources.

1. In the navigation pane, click **Administration > Single Sign On**.
2. Click **Identity Provider**.
3. Click **Add** and select **AD over LDAP** or **OpenLDAP**.

Connect Identity Provider

> ✔ Overview Connecting your identity provider

2. AD over LDAP selected

Select the identity provider you want to assign to the vCenter server.

Select Identity Provider
Embedded ▼

Embedded supports LDAP and Open LDAP. Add the identity source below.

Select Identity Source
AD over LDAP ▼

NEXT

3. Server Settings Active Directory over LDAP settings

4. Review Review information and Submit

4. Click **Next**.
5. Enter the server settings and click **Next**.

Table 194: Active Directory over LDAP and OpenLDAP Server Settings

Option	Description
Identity Source Name	Name of the identity source.
Base Distinguished Name for Users	Base Distinguished Name for users. Enter the DN from which to start user searches. For example, <code>cn=Users,dc=myCorp,dc=com</code> .
Base Distinguished Name for Groups	The Base Distinguished Name for groups. Enter the DN from which to start group searches. For example, <code>cn=Groups,dc=myCorp,dc=com</code> .
Domain Name	The FQDN of the domain.
Domain Alias	For Active Directory identity sources, the domain's NetBIOS name. Add the NetBIOS name of the Active Directory domain as an alias of the identity source if you are using SSPI authentications. For OpenLDAP identity sources, the domain name in capital letters is added if you do not specify an alias.
User Name	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups. The ID can be in any of these formats: <ul style="list-style-type: none"> • UPN (user@domain.com)

Table continued on next page

Continued from previous page

Option	Description
	<ul style="list-style-type: none"> • NetBIOS (DOMAIN\user) • DN (cn=user,cn=Users,dc=domain,dc=com) The user name must be fully-qualified. An entry of "user" does not work.
Password	Password of the user who is specified by Username .
Primary Server URL	Primary domain controller LDAP server for the domain. You can use either the host name or the IP address. Use the format <code>ldap://hostname_or_IPaddress:port</code> or <code>ldaps://hostname_or_IPaddress:port</code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use <code>ldaps://</code> in the primary or the secondary LDAP URL.
Secondary Server URL	Address of a secondary domain controller LDAP server that is used for failover. You can use either the host name or the IP address.
Certificates (for LDAPS)	If you want to use LDAPS with your Active Directory LDAP Server or OpenLDAP Server identity source, click Browse to select a certificate. To export the root CA certificate from Active Directory, consult the Microsoft documentation.

6. Review the information and click **Submit**.

After you successfully add an identity source, you can add users and groups from the domain. See [Add a User or Group to](#) .

Configure Microsoft ADFS as the Identity Provider in the SDDC Manager UI

You can configure VMware Cloud Foundation to use Microsoft ADFS as an external identity provider, instead of using vCenter Single Sign-On. In this configuration, the external identity provider interacts with the identity source on behalf of vCenter Server.

Microsoft Active Directory Federation Services (ADFS) requirements:

- Microsoft ADFS for Windows Server 2016 or later must already be deployed.
- Microsoft ADFS must be connected to Active Directory.
- You have created a vCenter Server administrators group in Microsoft ADFS that contains the users you want to grant vCenter Server administrator privileges to.

For more information about configuring Microsoft ADFS, see the Microsoft documentation.

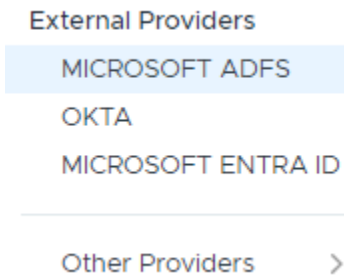
vCenter Server and other requirements:

- vSphere 7.0 or later

- vCenter Server must be able to connect to the Microsoft ADFS discovery endpoint, and the authorization, token, logout, JWKS, and any other endpoints advertised in the discovery endpoint metadata.

You can only add one external identity provider to VMware Cloud Foundation.

1. Log in to the SDDC Manager UI as a user with the ADMIN role
2. In the navigation pane, click **Administration** › **Single Sign On**.
3. Click **Identity Provider**.
4. Click **Change Identity Provider** and select **Microsoft ADFS**.



5. Click **Next**.
6. Select the checkbox to confirm the prerequisites and click **Next**.
7. If your Microsoft ADFS server certificate is signed by a publicly trusted Certificate Authority, click **Next**. If you are using a self-signed certificate, add the Microsoft ADFS root CA certificate added to the Trusted Root Certificates Store.
 - a) Click **Browse**.
 - b) Navigate to the certificate and click **Open**.
 - c) Click **Next**.
8. Copy the redirect URIs.

You will need them when you create the Microsoft ADFS Application Group in the next step.

9. Create an OpenID Connect configuration in Microsoft ADFS.

To establish a relying party trust between vCenter Server and an identity provider, you must establish the identifying information and a shared secret between them. In Microsoft ADFS, you do so by creating an OpenID Connect configuration known as an Application Group, which consists of a Server application and a Web API. The two components specify the information that vCenter Server uses to trust and communicate with the Microsoft ADFS server. To enable OpenID Connect in Microsoft ADFS, see the VMware knowledge base article at <https://kb.vmware.com/s/article/78029>.

Note the following when you create the Microsoft ADFS Application Group.

- You need the two Redirect URIs from the previous step.
 - Copy the following information to a file or write it down for use when configuring the identity provider in the next step.
 - Client Identifier
 - Shared Secret
 - OpenID address of the Microsoft ADFS server
10. Enter the Application Group information and click **Next**.
Use the information you gathered in the previous step and enter the:
 - Client Identifier
 - Shared Secret

- OpenID address of the Microsoft ADFS server

11. Enter user and group information for the Active Directory over LDAP connection to search for users and groups.

vCenter Server derives the AD domain to use for authorization and permissions from the Base Distinguished Name for users. You can add permissions on vSphere objects only for users and groups from this AD domain. Users or groups from AD child domains or other domains in the AD forest are not supported by vCenter Server Identity Provider Federation.

Option	Description
Base Distinguished Name for Users	Base Distinguished Name for users.
Base Distinguished Name for Groups	The base Distinguished Name for groups.
User Name	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Password	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Primary Server URL	Primary domain controller LDAP server for the domain. Use the format <code>ldap://hostname:port</code> or <code>ldaps://hostname:port</code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use <code>ldaps://</code> in the primary or secondary LDAP URL.
Secondary Server URL	Address of a secondary domain controller LDAP server that is used for failover.
Certificates (for LDAPS)	If you want to use LDAPS, click Browse to select a certificate.

12. Review the information and click **Submit**.

After you successfully add Microsoft ADFS as an external identity provider, you can add users and groups to VMware Cloud Foundation. See [Add a User or Group to](#) .


Add a User or Group to VMware Cloud Foundation

You can add users or groups so that they can access the SDDC Manager UI and VMware Cloud Foundation API.

Only a user with the ADMIN role can perform this task.

SDDC Manager UI displays user and group information based on the configured identity provider and identity sources. See [Configuring the Identity Provider for VMware Cloud Foundation](#).

1. In the navigation pane, click **Administration** > **Single Sign On**.
2. Click **Users and Groups** and then click **+ User or Group**.



3. Select one or more users or group by clicking the check box next to the user or group.

You can either search for a user or group by name, or filter by user type or domain.

4. Select a Role for each user and group.

Role	Description
ADMIN	This role has access to all the functionality of the UI and API.
OPERATOR	This role cannot access user management, password management, or backup configuration settings.
VIEWER	This role can only view the SDDC Manager. User management and password management are hidden from this role.

5. Scroll down to the bottom of the page and click **Add**.

Remove a User or Group

You can remove a user or group, for example when an employee leaves the company. The removed user or group will not be able to log in to the SDDC Manager UI.

Only a user with the ADMIN role can perform this task.

1. In the navigation pane, click **Administration > Single Sign On**.
2. Click the vertical ellipsis (three dots) next to a user or group name and click **Remove**.
3. Click **Delete**.

Create a Local Account

A local account is used to access VMware Cloud Foundation APIs when the management vCenter Server is down. If you upgraded from a previous release or didn't configure the account when deploying using the API, you can set a password using VMware Cloud Foundation APIs.

1. Log in to the SDDC Manager UI as a user with the ADMIN role.
For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).
2. In the navigation pane, click **Developer Center > API Explorer**.
3. To verify if the local account is configured, perform the following tasks:
 - a) Expand **APIs for managing Users**.
 - b) Expand `GET /v1/users/local/admin` and click **EXECUTE**.
 - c) In the Response, click `LocalUser (admin@local)`.

Response

```
LocalUser (admin@local) [🔗] ⬇ {  
  "isConfigured":  
    Flag indicating whether or not local  
    account is configured  
  "true",  
  "name":  
    The name of the user  
  "admin@local",  
  "role":  
    The role of the user  
  RoleReference (fa16f14b-9679-bbfc-06ed-47245405542c) [🔗] ⬇ { ... }  
  "type":  
    The type of the user  
  "USER",  
}
```

You can also download the response by clicking the download icon to the right of `LocalUser (admin@local)`.

4. If the local account is not configured, perform the following tasks to configure the local account:
 - a) Expand `PATCH /v1/users/local/admin`.
 - b) Enter a password for the local account and click **EXECUTE**.

PATCH /v1/users/local/admin Update password for local account

▼ Description
Update the password for local account only if the old password is correct, or if user configures the local account for the first time

> Response Types

▼ Try it out

Parameter	Value	Type	Description/Data Type
localUserPassword (required)	<pre>1 { 2 "newPassword": "Vmware123!Vmware123!" 3 }</pre>	Body	Local user password details LocalAccountPasswordInfo { newPassword: (string, required) The new password of the local account oldPassword: (string) The old password of the local account }

EXECUTE COPY JSON DOWNLOAD

Password requirements are described below:

- Minimum length: 12
- Maximum length: 127
- At least one lowercase letter, one uppercase letter, a number, and one of the following special characters ! % @ \$ ^ # ? *
- A character cannot be repeated more than three times consecutively
- Must not include three of the same consecutive characters

NOTE

You must remember the password that you created because it cannot be retrieved. Local account passwords are used in password rotation.

Create an Automation Account

Automation accounts are used to access VMware Cloud Foundation APIs in automation scripts.

1. Log in to the SDDC Manager UI as a user with the ADMIN role.
For more about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).
2. In the navigation pane, click **Developer Center** > **API Explorer**.
3. Get the ID for the ADMIN role.
 - a) Expand **APIs for managing Users**.
 - b) Expand **GET /v1/roles** and click **Execute**.
 - c) In the Response, click **PageOfRole** and **Role (ADMIN)**.
 - d) Copy the ID for the ADMIN role.

Response

```

PageOfRole [ ] ↓ {
  "elements":
    The list of elements included in this page
    [
      Role (ADMIN) [ ] ↓ {
        "description":
          The description of the role
          "Administrator",
        "id":
          The ID of the role
          "317cb292-802f-ca6a-e57e-3ac2b707fe34",
        "name":
          The name of the role
          "ADMIN",
      },
    ],
  },

```

4. Create a service account with the ADMIN role and get the service account's API key.
 - a) Expand **POST /v1/users** and click **User**.
 - b) Replace the Value with:

```

[
  {
    "name": "service_account",
    "type": "SERVICE",
    "role":
      {
        "id": "317cb292-802f-ca6a-e57e-3ac2b707fe34"
      }
  }
]

```

Paste the ADMIN role ID from step 3.

POST /v1/users Add users

Description
Add list of users

Response Types
Try it out

Parameter	Value	Type	Description/Data Type
users (required)	<pre> 1 [2 { 3 "name": "service", 4 "role": { 5 "id": "317cb292-802f-ca6a-e57e-3ac2b70", 6 }, 7 "type": "SERVICE" 8 } 9] </pre>	Body	User data collection [User{ ... }]

EXECUTE COPY JSON DOWNLOAD

- Click **Execute**.
- In the Response, click `PageOfUser` and `User (service_account)`.
- Copy the API key for the service account.

Response

`PageOfUser` {

"elements":
The list of elements included in this page
[

`User (service_account)` {

"apiKey":
The API key of the user
"qsfqnYgyxXQ892Jk9OHXyuEMgE3SgfTS",

- Use the service account's API key to generate an access token.

- a) Expand **APIs for managing access and refresh tokens**.
- b) Expand **POST /v1/tokens**.
- c) Click **TokenCreationSpec**.
- d) Replace Value with:

```
{
  "apiKey": "qsfcqnYgyxXQ892Jk90HXyuEMgE3SgfTS"
}
```

Paste the service account's API key from step 4.

APIs for managing access and refresh token

POST /v1/tokens Create Token Pair

Description
Creates access token and refresh token for user access

Response Types

Try it out

Parameter	Value	Type	Description/Data Type
tokenCreationSpec (required)	<pre>{ 1: { 2: "apiKey": "qsfcqnYgyxXQ892Jk90HXyuEMgE3SgfTS" 3: } }</pre>	Body	tokenCreationSpec TokenCreationSpec{ ... }

EXECUTE COPY JSON DOWNLOAD

- e) Click **Execute**.
- f) In the Response, click **TokenPair** and **RefreshToken** and save the access and refresh tokens.

Response

TokenPair {

"accessToken":
Bearer token that can be used to make public API calls
"eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXZWQ9MTUzZmZlLTQwODMtYjU3YjY1fH0DYxOTUwODY0YTQlCjYyXQ0jE1ODU1ODgyNDAsInN1..."

"refreshToken":
Refresh token that can be used to request new access token

RefreshToken (33f88c60-862e-4a38-8e8e-6479c4cd9f33) {

"id":
Refresh token id that can be used to request new access token
"33f88c60-862e-4a38-8e8e-6479c4cd9f33",

}

}

Managing Passwords in VMware Cloud Foundation

For security reasons, you can change passwords for the accounts that are used by your SDDC Manager instance. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

NOTE

For information about password policy design including the details and justification for the configuration of password expiration, complexity, and account lockout policies, see the [Information Security and Access](#) in the *Identity and Access Management for VMware Cloud Foundation* validated solution. For step-by-step instructions on configuring password policies, see [Password Policy Configuration for VMware Cloud Foundation](#).

You entered passwords for your VMware Cloud Foundation system as part of the bring-up procedure. You can rotate and update some of these passwords using the password management functionality in the SDDC Manager UI, including:

- Accounts used for service consoles, such as the ESXi root account.
- The single sign-on administrator account(s).

NOTE

SDDC Manager manages passwords for all SSO administrator accounts, even if you created isolated VI workload domains that use different SSO domains than the management domain.

- The default administrative user account used by virtual appliances.
- Service accounts that are automatically generated during bring-up, host commissioning, and workload creation. Service accounts have a limited set of privileges and are created for communication between products. Passwords for service accounts are randomly generated by SDDC Manager. You cannot manually set a password for service accounts. To update the credentials of service accounts, you can rotate the passwords.

To provide optimal security and proactively prevent any passwords from expiring, you must rotate passwords every 80 days.

NOTE

Do not change the passwords for system accounts and the `administrator@vsphere.local` account outside SDDC Manager. This can break your VMware Cloud Foundation system.

You can also use the VMware Cloud Foundation API to look up and manage credentials. In the SDDC Manager UI, click **Developer Center** › **API Explorer** and browse to the APIs for managing credentials.

Starting with VMware Cloud Foundation 5.2.1, you can also manage passwords using the vSphere Client.

Password Expiration Notifications

The SDDC Manager UI provides a banner notification for any passwords managed by VMware Cloud Foundation that are expiring within the next 14 days. For example:



5 passwords will expire within 14 days. Visit Password Management page to take action.

MANAGE



You can also click **Security** › **Password Management** in the navigation pane to view password expiration information.

NOTE

Password expiration information is not available for vCenter Server SSO service accounts.

Expired passwords will display a status of `Disconnected`. For example:

ESXi vCenter PSC NSXT Manager NSXT Edge VRS LCM Backup						
ROTATE NOW		SCHEDULE ROTATION ▾				
<input type="checkbox"/>	User Name	FQDN	Domain	IP Address	Status	
<input type="checkbox"/>	root	esx-7.vrack.vsphere.local	MGMT	10.0.0.101	Disconnected on 1/21/30, 12:00 AM.	
<input type="checkbox"/>	root	esx-45.vrack.vsphere.local	MGMT	10.0.0.102	Disconnected on 1/19/30, 8:23 AM.	
<input type="checkbox"/>	root	esx-qrs.vrack.vsphere.local	MGMT	10.0.0.103	Disconnected on 1/4/30, 8:36 AM.	

For an expired password, you must update the password outside of VMware Cloud Foundation and then remediate the password using the SDDC Manager UI or the VMware Cloud Foundation API. See [Remediate Passwords](#).

NOTE

Password expiration information in the SDDC Manager UI is updated once a day. To get real-time information, use the VMware Cloud Foundation API.

Rotate Passwords

As a security measure, you can rotate passwords for the components in your VMware Cloud Foundation instance. The process of password rotation generates randomized passwords for the selected accounts. You can rotate passwords manually or set up auto-rotation for accounts managed by SDDC Manager.

- Verify that there are no currently failed workflows in SDDC Manager. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the brief time period that the password rotation process is running. It is recommended that you schedule password rotation for a time when you expect to have no running workflows.
- Only a user with the ADMIN role can perform this task.

You can rotate passwords for the following accounts.

- ESXi

NOTE

Auto-rotate is not supported for ESXi.

- vCenter Server
By default, the vCenter Server root password expires after 90 days.

NOTE

Auto-rotate is automatically enabled for vCenter Server service accounts. It may take up to 24 hours to configure the service account auto-rotate policy for a newly deployed vCenter Server.

- vSphere Single-Sign On (PSC)
- NSX Edge nodes
- NSX Manager
- VMware Avi Load Balancer (formerly known as NSX Advanced Load Balancer)
- VMware Aria Suite Lifecycle
- VMware Aria Operations for Logs
- VMware Aria Operations
- VMware Aria Automation
- Workspace ONE Access

NOTE

For Workspace ONE Access passwords, the password rotation method varies depending on the user account. See the table below for details.

- SDDC Manager`backup` user

Table 195: Password Rotation Details for Workspace ONE Access User Accounts

Workspace ONE Access User Account	VMware Aria Suite Lifecycle Locker Entry	Password Rotation Method	Password Rotation Scope
admin (443)	xint-wsa-admin	SDDC Manager Password Rotation	Application
admin (8443)	xint-wsa-admin	VMware Aria Suite Lifecycle Global Environment	Per node
configadmin (443)	xint-wsa-configadmin	<ol style="list-style-type: none"> 1. Reset the configadmin user password in Workspace ONE Access via the email reset link. 2. Create a new credential object in VMware Aria Suite Lifecycle Locker to match the new password. 3. Update the credential object referenced by <code>globalEnvironment</code> in VMware Aria Suite Lifecycle locker to the new credential object. 	Application
sshuser	global-env-admin	VMware Aria Suite Lifecycle Global Environment	Per node
root (ssh)	xint-wsa-root	SDDC Manager Password Rotation	Per node

The default password policy for rotated passwords requires:

- 20 characters in length
- At least one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ *
- No more than two of the same characters consecutively

If you changed the vCenter Server password length using the vSphere Client or the ESXi password length using the VMware Host Client, rotating the password for those components from SDDC Manager generates a password that complies with the password length that you specified.

To update the SDDC Manager root, super user, and API passwords, see [Updating SDDC Manager Passwords](#).

1. In the navigation pane, click **Security > Password Management**.
2. Select one or more accounts and click one of the following operation.
 - **Rotate Now**
 - **Schedule Rotation**
You can set the password rotation interval (30 days, 60 days, or 90 days). You can also deactivate the schedule.

NOTE

The **Schedule Rotation** option is not available for ESXi.

NOTE

Auto-rotate schedule is configured to run at midnight on the scheduled date. If auto-rotate could not start due to any technical issue, there is a provision to auto-retry every hour till start of the next day. In case schedule rotation is missed due to technical issues the UI displays a global notification with failed task status. The status of the schedule rotation can also be checked on the Tasks panel.

A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status for the password rotation operation. To view sub-tasks, click the task name. As each of these tasks is run, the status is updated. If the task fails, you can click **Retry**.

Password rotation is complete when all sub-tasks are completed successfully.

Manually Update Passwords

You can manually change the password for a selected account. Unlike password rotation, which generates a randomized password, you provide the new password.

- Verify that there are no currently failed workflows in your VMware Cloud Foundation system. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the manual password update.
- Only a user with the ADMIN role can perform this task. For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).

You can update only one password at a time.

Although individual VMware Cloud Foundation accounts support different password requirements, it is recommended that you set passwords following a common set of requirements across all accounts:

- Minimum length: 15
- Maximum length: 20
- At least one lowercase letter, one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ *
- Must NOT include:
 - A dictionary word
 - A palindrome
 - More than four monotonic character sequences
 - Three of the same consecutive characters

1. From the navigation pane, select **Security > Password Management**.
2. Select the account whose password you want to update, click the vertical ellipsis (three dots), and click **Update Password**.
3. Enter and confirm the new password.
4. Click **Update**.

A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status of the password update operation. To view sub-tasks, click the task name.

Password update is complete when all sub-tasks are completed successfully.

Remediate Passwords

When an error occurs, for example after a password expires, you must manually reset the password in the component product. After you reset the password in a component, you must remediate the password in SDDC Manager to update the password in the SDDC Manager database and the dependent VMware Cloud Foundation workflows.

- Verify that VMware Cloud Foundation system contain no failed workflows. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no workflows are running or are scheduled to run while you remediate the password.
- Only a user with the ADMIN role can perform this task. For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).

To resolve any errors that might have occurred during password rotation or update, you must use password remediation. Password remediation syncs the password of the account stored in the SDDC Manager with the updated password in the component.

NOTE

You can remediate the password for only one account at a time.

Although the individual VMware Cloud Foundation components support different password requirements, you must set passwords following a common set of requirements across all components.

For information on updating passwords manually, see [Manually Update Passwords](#).

1. From the navigation pane, select **Security** › **Password Management**.
2. Select the account whose password you want to remediate, click the vertical ellipsis (three dots), and click **Remediate Password**.

The Remediate Password dialog box appears. This dialog box displays the entity name, account type, credential type, and user name, in case you must confirm you have selected the correct account.

3. Enter and confirm the password that was set manually on the component.
4. Click **Remediate**.

A message appears at the top of the page showing the progress of the operation. The Task panel also shows detailed status of the password remediation operation. To view subtasks, you can click the task name.

Password remediation is complete when all sub-tasks are completed successfully.

Look Up Account Credentials

To look up the account credentials for the built-in accounts that are managed and rotated by SDDC Manager, you can log in to the SDDC Manager appliance using any SDDC Manager account credentials.

Only a user with the ADMIN role can perform this task.

1. SSH in to the SDDC Manager appliance using the `vcf` user account.
2. Change to the `/usr/bin` directory.

NOTE

Although the password management CLI commands are located in `/usr/bin`, you can run them from any directory.

3. Obtain the account credentials list by typing the command:

```
lookup_passwords
```

You must enter the user name and password for a user with the ADMIN role.

NOTE

Accounts with type `USER` and `SYSTEM` will be displayed.

4. Save the command output to a secure location with encryption so that you can access it later and use it to log in to the accounts as needed.

Updating SDDC Manager Passwords

The process for updating SDDC Manager passwords varies, depending on which account you are updating.

Update SDDC Manager Root and Super User Passwords

For security reasons, you can change passwords for the SDDC Manager root (`root`) and super user (`vcf`) accounts. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

The SDDC Manager `root` password expires after 90 days.

1. SSH in to the SDDC Manager VM using the `vcf` user account.
2. Enter `su` to switch to the root user.
3. Enter one of the following commands:

<code>passwd vcf</code>	To change the super user password.
<code>passwd root</code>	To change the root password.

4. Enter and retype the new password. For example:

```
root@sddc-manager [ /home/vcf ]# passwd vcf
New password:
Retype new password:
passwd: password updated successfully
```

The password is updated.

Update SDDC Manager Local Account Password

The SDDC Manager local account is used to access VMware Cloud Foundation APIs when the management vCenter Server is down. For security reasons, you should periodically update the password for this account.

Password requirements for the SDDC Manager local account:

- At least 15 characters
- No more than 127 characters
- At least one lowercase letter
- At least one uppercase letter
- At least one digit
- At least one special character, such as `@ ! # $ % ^` or `?`
- A character cannot be repeated more than 3 times consecutively

1. Log in to the SDDC Manager UI as a user with the ADMIN role.

For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).

2. Click **Developer Center > API Explorer**.
3. Expand **APIs for managing Users**.
4. Expand **PATCH /v1/users/local/admin**.
5. In the **Description/Data Type** column, click **LocalAccountPasswordInfo{...}**.
6. In the **Value** box, type the new and old passwords and click **Execute**.

- Click **Continue** to confirm.

A response of `Status: 204, No Content` indicates that the password was successfully updated.

Update Expired SDDC Manager Root Password

This section describes the procedure for updating an expired password for the SDDC Manager root (`root`) user.

The password must meet the following requirements:

- Minimum length 15 characters
 - Must include:
 - mix of uppercase and lowercase letters
 - a number
 - a special character, such as `@ ! # $ % ^` or `?`
 - Must not include:
 - `* { } [] () / \ ' " ` ~ , ; : . < >`
 - A dictionary word (for example, `VMware1!`)
1. In a web browser, log in to the management domain vCenter Server using the vSphere Client (`https://<vcenter_server_fqdn>/ui`).
 2. In the VMs and Templates inventory, expand the management domain vCenter Server and the management virtual machines folder.
 3. Right-click the SDDC Manager virtual machine, and select **Open Remote Console**.
 4. Click within the console window and press **Enter** on the Login menu item.
 5. Type **root** as the user name and enter the current password for the root user.
 6. Type `passwd root`.
 7. When prompted for a new password, enter a different password than the previous one and click **Enter**.

Backup and Restore of VMware Cloud Foundation

Regular backups of the management components ensures that you can keep your environment operational if a data loss or failure occurs.

You implement scheduled backups to prepare for:

- A critical failure of a management component.
- An upgrade of a management component.
- A certificate update of a management component.

In addition, as a best practice, you can take on-demand manual backups in the following use cases:

- After a successful recovery operation.
- After resolving asynchronously reported errors in SDDC components.
- After resolving an incomplete workflow in SDDC Manager.
- After noting the failure of a scheduled backup of an SDDC component.
- Immediately before performing a system upgrade.

You can backup and restore SDDC Manager with a file-based or an image-based solution.

For a file-based backup of the SDDC Manager appliance, you can configure a backup schedule and enable task-based (state-change driven) backups. When task-based backups are enabled, a backup is triggered after each SDDC Manager task such as workload domain and host operations or password rotation. You can also define a backup retention policy to comply with your company's retention policy.

By default, NSX Manager file-based backups are taken and stored on an SFTP server that is built into the SDDC Manager appliance. It is recommended that you configure an external SFTP server as a backup location for the following reasons:

- An external SFTP server is a prerequisite for restoring SDDC Manager file-based backups.
- Using an external SFTP server provides better protection against failures because it decouples NSX backups from SDDC Manager.

By default, VMware Cloud Foundation does not configure any file-based or image-based backup protection for the vCenter Servers.

This section of the documentation provides instructions on backing up and restoring SDDC Manager, NSX, and vCenter Server.

Reconfigure SFTP Backups for SDDC Manager and NSX Manager

By default, backups of SDDC Manager and NSX Manager are stored in the SDDC Manager appliance. Change the destination of the backups to an external SFTP server.

- Only a user with the ADMIN role can perform this task. See [Managing Users and Groups in VMware Cloud Foundation](#).
- The external SFTP server must support a 256-bit length ECDSA SSH public key.
- The external SFTP server must support a 2048-bit length RSA SSH public key
- You will need the SHA256 fingerprint of RSA key of the SFTP server.
- Host Key algorithms: At least one of rsa-sha2-512 or rsa-sha2-256 and one of ecdsa-sha2-nistp256, ecdsa-sha2-nistp384, or ecdsa-sha2-nistp521.
- Additional pre-requisites when FIPS Security Mode is enabled on SDDC Manager:

Algorithms and Ciphers	Required when FIPS Security Mode is Enabled
Kex Algorithms	At least one of: <ul style="list-style-type: none"> • diffie-hellman-group-exchange-sha256 • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521
Message Authentication Key (MAC) Algorithms	hmac-sha2-256
Ciphers	At least one of: <ul style="list-style-type: none"> • TLS_DHE_DSS_WITH_AES_128_CBC_SHA256 • TLS_DHE_DSS_WITH_AES_256_CBC_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 • TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 • TLS_RSA_WITH_AES_128_CBC_SHA256 • TLS_RSA_WITH_AES_256_CBC_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 • TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 • TLS_AES_128_GCM_SHA256 • TLS_AES_256_GCM_SHA384

NOTE

SHA1 algorithms are not supported.

1. In the navigation pane, click **Administration** > **Backup**.
2. On the **Backup** page, click the **Site Settings** tab and then click **Register External**.
3. On the **Backup** page, enter the settings and click **Save**.

To obtain the SSH Fingerprint of the target system to verify, connect to the SDDC Manager Appliance over ssh and run the following command:

```
ssh-keygen -lf <(ssh-keyscan -p 22 -t rsa sftp_server_fqdn 2> /dev/null) | cut -d' ' -f2
```

Setting	Value
Host FQDN or IP	The FQDN or IP Address of the SFTP server.
Port	22
Transfer Protocol	SFTP
Username	A service account with privileges to the SFTP server. For example: <code>svc-vcf-bck</code> .
Password	The password for the username provided.
Backup Directory	The directory on the SFTP server where backups are saved. For example: <code>/backups/</code> .
SSH Fingerprint	The SSH Fingerprint is automatically retrieved from the SFTP server, verify the SSH Fingerprint.
Confirm Fingerprint	Selected
Encryption Passphrase	The encryption passphrase used to encrypt the backup data. NOTE The encryption passphrase should be stored safely as it is required during the restore process.

4. In the **Confirm your changes to backup settings** dialog box, click **Confirm**.

File-Based Backups for SDDC Manager and vCenter Server

You can use the native file-based backup capabilities of SDDC Manager, vCenter Server, and NSX Manager. The NSX Manager backup is configured by SDDC Manager during the bring-up process. You configure the file-based backup jobs for SDDC Manager and vCenter Server.

Verify that you have an SFTP server on the network to serve as a target of the file-based backups.

To ensure that all management components are backed up correctly, you must create a series of backup jobs that capture the state of a set of related components at a common point in time. With some components, simultaneous backups of the component nodes ensure that you can restore the component a state where the nodes are logically consistent with each other and eliminate the necessity for further logical integrity remediation of the component.

Table 196: File-Based Backup Jobs

Component	Recommended Frequency	Recommended Retention	Notes
SDDC Manager	Daily	7 days	You must configure the backup jobs for the SDDC Manager instance and all vCenter Server instances in the vCenter Single Sign-On domain to start within the same 5-minute window.
vCenter Server	Daily	7 days	
vSphere Distributed Switch	On-demand	Retain last 3 configurations.	-
NSX Manager	Hourly	7 days	Configured by SDDC Manager during the bring-up process.

NOTE

- You must monitor the space utilization on the SFTP server to ensure that you have sufficient storage space to accommodate all backups taken within the retention period.
- Do not make any changes to the `/opt/vmware/vcf` directory on the SDDC Manager VM. If this directory contains any large files, backups may fail.

Back Up SDDC Manager

You configure file-based daily backups of the SDDC Manager instances using the SDDC Manager administration interface.

Only a user with the **Admin** role can perform this task.

1. In the navigation pane, click **Administration > Backup**.
2. On the **Backup** page, click the **SDDC Manager Configurations** tab.
3. Under **Backup Schedule**, click **Edit**.
4. On the **Backup Schedule** page, enter the settings and click **Save**.

Setting	Value
Automatic Backup	Enabled
Backup Frequency	Weekly
Days of the Week	All selected
Schedule Time	04:02 AM
Take Backup on State Change	Enabled
Retain Last Backups	7
Retain Hourly Backups for Days	1
Retain Daily Backups for Days	7

5. To verify the backup, click **Backup Now**.

The status and the start time of the backup is displayed on the UI. You have set the SDDC Manager backup schedule to run daily at 04:02 AM and after each change of state.

If the backup is unsuccessful, verify if the SFTP server is available and able to provide its SSH fingerprint:

- SSH to the SDDC Manager appliance run the following command as the root user:

```
sftp username@IP of sftp server
```

Enter the SFTP user password when prompted. The following message indicates a successful connection:

```
Connected to username@IP of sftp server.
```

- To check that the SFTP server SSH fingerprint is available, run:

```
ssh-keygen -lf <(ssh-keyscan -t ssh-rsa -p port_numberIP of sftp server 2>/dev/null)
```

Configure a Backup Schedule for vCenter Server

You configure file-based daily backups of the vCenter Server instances by using the vCenter Server Management Interface of each vCenter Server instance.

1. In a web browser, log in to the vCenter Server Management Interface (<https://appliance-IP-address-or-FQDN:5480>).
2. In the left navigation pane, click **Backup**.
3. In the **Backup schedule** pane, click **Configure**.
4. In the **Create backup schedule** dialog box, enter these values and click **Create**.

Setting		Value
Backup location		Enter the backup location from SFTP server. For example: <code>sftp://172.16.11.60/backups/</code>
Backup server credentials	User name	A service account with privileges to the SFTP server. For example: <code>svc-vcf-bck</code> .
	Password	Enter the password for the username provided.
Schedule		Daily 11:00 PM
Encrypt backup	Encryption password	<code>encryption_password</code>
	Confirm password	<code>encryption_password</code>
Number of backups to retain		Retain last 7 backups
Data	Stats, events, and tasks	Selected
	Inventory and configuration	Selected

The backup schedule information appears in the **Backup schedule** pane.

5. Repeat the procedure for the other vCenter Server instances.

Any complete and in-progress backup appears in the **Activity** pane.

Manually Back Up vCenter Server

Before you upgrade a vCenter Server instance, you should use the vCenter Server Management Interface to manually back it up.

- In the vSphere Client, for each vSphere cluster that is managed by the vCenter Server, note the current vSphere DRS Automation Level setting and then change the setting to **Manual**. After the vCenter Server upgrade is complete,

you can change the vSphere DRS Automation Level setting back to its original value. See [KB 87631](#) for information about using VMware PowerCLI to change the vSphere DRS Automation Level.

- Ensure that there are not any active vMotion tasks.
1. In a web browser, log in to the vCenter Server Management Interface (<https://appliance-IP-address-or-FQDN:5480>).
 2. In the left navigation pane, click **Backup**.
 3. Click **Backup Now**.
 4. If you already have a backup schedule set up, select **Use backup location and user name from backup schedule** and click **Start**.
 5. If you do not already have a backup schedule, enter the following information and click **Start**.

Setting		Value
Backup location		Enter the backup location from SFTP server. For example: <code>sftp://172.16.11.60/backups/</code>
Backup server credentials	User name	A service account with privileges to the SFTP server. For example: <code>svc-vcf-bck</code> .
	Password	Enter the password for the username provided.
Encrypt backup	Encryption password	<code>encryption_password</code>
	Confirm password	<code>encryption_password</code>
Data	Stats, events, and tasks	Selected
	Inventory and configuration	Selected

In order to restore vCenter Server, you will need the VMware vCenter Server Appliance ISO file that matches the version you backed up.

- Identify the required vCenter Server version. In the vCenter Server Management Interface, click **Summary** in the left navigation pane to see the vCenter Server version and build number.
- Download the VMware vCenter Server Appliance ISO file for that version from the Broadcom Support Portal.

Export the Configuration of the vSphere Distributed Switches

The vCenter Server backup includes the configuration of the entire vCenter Server instance. To have a backup only of the vSphere Distributed Switch and distributed port group configurations, you export a configuration file that includes the validated network configurations. If you want to recover only the vSphere Distributed Switch, you can import this configuration file to the vCenter Server instance.

You can use the exported file to create multiple copies of the vSphere Distributed Switch configuration on an existing deployment, or overwrite the settings of existing vSphere Distributed Switch instances and port groups. You must backup the configuration of a vSphere Distributed Switch immediately after each change in configuration of that switch.

1. In a web browser, log in to vCenter Server by using the vSphere Client.
2. Select **Menu > Networking**.
3. In the inventory expand **vCenter Server > Datacenter**.

4. Expand the **Management Networks** folder, right-click the distributed switch, and select **Settings > Export configuration**.
5. In the **Export configuration** dialog box, select **Distributed switch and all port groups**.
6. In the **Description** text box enter the date and time of export, and click **OK**.
7. Copy the backup zip file to a secure location from where you can retrieve the file and use it if a failure of the appliance occurs.
8. Repeat the procedure for the other vSphere Distributed Switches.

File-Based Restore for SDDC Manager, vCenter Server, and NSX

When SDDC Manager, vCenter Server, or NSX Manager in the SDDC fails, you can restore the component to a fully operational state by using its file-based backup. When an NSX Edge node fails, you redeploy the node from the NSX Manager instance.

Use this guidance as appropriate based on the exact nature of the failure encountered within your environment. Sometimes, you can recover localized logical failures by restoring individual components. In more severe cases, such as a complete and irretrievable hardware failure, to restore the operational status of your SDDC, you must perform a complex set of manual deployments and restore sequences. In failure scenarios where there is a risk of data loss, there has already been data loss or where it involves a catastrophic failure, contact Broadcom Support to review your recovery plan before taking any steps to remediate the situation.

Restore SDDC Manager

If SDDC Manager fails, you can restore it from its file-based backup.

- Power off and rename the failed SDDC Manager instance.
- Verify that you have a valid file-based backup of the failed SDDC Manager instance.

To be valid, the backup must be of the same version as the version of the SDDC Manager appliance on which you plan to restore the instance.

- Verify that you have the SFTP server details:
 - SFTP Server IP
 - SFTP Server Username
 - SFTP Server Password
 - Encryption Password

After a successful recovery, securely delete the decrypted backup files.

Prepare for Restoring SDDC Manager

Before restoring SDDC Manager, you must download and decrypt the encrypted backup file from the SFTP server.

Verify that your host machine with access to the SDDC has OpenSSL installed.

NOTE

The procedures have been written based on the host machine being a Linux-based operating system.

The backup file contains sensitive data about your VMware Cloud Foundation instance, including passwords in plain text. As a best practice, you must control access to the decrypted files and securely delete them after you complete the restore operation.

1. Identify the backup file for the restore and download it from the SFTP server to your host machine.
2. On your host machine, open a terminal and run the following command to extract the content of the backup file.

```
OPENSSL_FIPS=1 openssl enc -d -aes-256-cbc -md sha256 -in filename-of-restore-file |
tar -xz
```

3. When prompted, enter the *encryption_password*.
4. In the extracted folder, locate and open the `metadata.json` file in a text editor.
5. Locate the `sddc_manager_ova_location` value and copy the URL.
6. In a web browser, paste the URL and download the OVA file.
7. In the extracted folder, locate and view the contents of the `security_password_vault.json` file.
8. Locate the `entityType BACKUP` value and record the backup password.

Restore SDDC Manager from a File-Based Backup

First, you deploy a new SDDC Manager appliance by using the OVA file that you downloaded during the preparation for the restore. After that, you restore the file-based backup on the newly deployed SDDC Manager appliance.

1. In a web browser, log in to management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** > **VMs and templates**.
3. In the inventory expand **vCenter Server** > **Datacenter**.
4. Right-click the management folder and select **Deploy OVF template**.
5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, browse to the location of the SDDC Manager OVA file, click **Open**, and click **Next**.
6. On the **Select a name and folder** page, in the **Virtual machine name** text box, enter a virtual machine name, and click **Next**.
7. On the **Select a compute resource** page, click **Next**.
8. On the **Review details** page, review the settings and click **Next**.
9. On the **License agreements** page, accept the license agreement and click **Next**.
10. On the **Select storage** page, select the vSAN datastore and click **Next**.

The datastore must match the `vsan_datastore` value in the `metadata.json` file that you downloaded during the preparation for the restore.

11. On the **Select networks** page, from the **Destination network** drop-down menu, select the management network distributed port group and click **Next**.

The distributed port group must match the `port_group` value in the `metadata.json` file that you downloaded during the preparation for the restore.

12. On the **Customize template** page, enter the following values and click **Next**.

Setting	Description
Enter root user password	You can use the original root user password or a new password.
Enter login (vcf) user password	You can use the original vcf user password or a new password.
Enter basic auth user password	You can use the original admin user password or a new password.

Table continued on next page

Continued from previous page

Setting	Description
Enter backup (backup) user password	The backup password that you saved during the preparation for the restore. This password can be changed later if desired.
Enter Local user password	You can use the original Local user password or a new password.
Hostname	The FQDN must match the <code>hostname</code> value in the <code>metadata.json</code> file that you downloaded during the preparation for the restore.
NTP sources	The NTP server details for the appliance.
Enable FIPs	Selected
Default gateway	The default gateway for the appliance.
Domain name	The domain name for the appliance.
Domain search path	The domain search path(s) for the appliance.
Domain name servers	The DNS servers for the appliance.
Network 1 IP address	The IP address for the appliance.
Network 1 netmask	The subnet mask for the appliance.

13. On the **Ready to complete** page, click **Finish** and wait for the process to complete.
14. When the SDDC Manager appliance deployment completes, expand the management folder.
15. Right-click the SDDC Manager appliance and select **Snapshots > Take Snapshot**.
16. Right-click the SDDC Manager appliance, select **Power > Power On**.
17. On the host machine, copy the encrypted backup file to the `/tmp` folder on the newly deployed SDDC Manager appliance by running the following command. When prompted, enter the `vcf_user_password`.

```
scp filename-of-restore-file vcf@sddc_manager_fqdn:/tmp/
```

18. On the host machine, obtain the authentication token from the SDDC Manager appliance in order to be able to execute the restore process by running the following command:

```
TOKEN=`curl https://<sddc_manager_fqdn>/v1/tokens -k -X POST -H "Content-Type: application/json" -d '{"username": "admin@local", "password": "<admin@local_password>"}' | awk -F "\"" '{ print $4}'`
```

19. On the host machine with access to the SDDC Manager, open a terminal and run the command to start the restore process.

```
curl https://<sddc_manager_fqdn>/v1/restores/tasks -k -X POST -H "Content-Type: application/json" -H "Authorization: Bearer $TOKEN" \
  -d '{
  "elements" : [ {
    "resourceType" : "SDDC_MANAGER"
  } ],
  "backupFile" : "<backup_file>",
  "encryption" : {
    "passphrase" : "<encryption_password>"
```

```
}
}'
```

The command output contains the ID of the restore task.

20. Record the ID of the restore task.
21. Monitor the restore task by using the following command until the status becomes `Successful`.

```
curl https://<sddc_manager_fqdn>/v1/restores/tasks/<restore_task_id> -k -X GET -H
"Content-Type: application/json" -H "Authorization: Bearer $TOKEN"
```

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Validate the Status of SDDC Manager

After a successful restore of SDDC Manager, you must validate its status. You run the health checks by using the `sos` tool.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client.
2. Run the health checks by using the `sos` tool.

```
sudo /opt/vmware/sddc-support/sos --health-check
```

3. When prompted, enter the `vcf_password`.
All tests show green when SDDC Manager is in healthy state.
4. Manually delete the snapshot created in [Restore SDDC Manager from a File-Based Backup](#).

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Restore vCenter Server

If a vCenter Server instance fails, you can restore it from its file-based backup.

- Power off and rename the failed vCenter Server instance.
- Verify that you have a valid file-based backup of the failed vCenter Server instance.

To be valid, the backup must be of the version of the vCenter Server Appliance on which you plan to restore the instance.

- Verify that you have the SFTP server details:
 - SFTP Server IP
 - SFTP Server Username
 - SFTP Server Password
 - Encryption Password

Prepare for Restoring vCenter Server

Before restoring a vCenter Server instance, you must retrieve the vCenter Server build number and deployment details, as well as vCenter Server and ESXi credentials from the SDDC Manager inventory.

SDDC Manager must be available.

Retrieve the vCenter Server Deployment Details

Before restoring a vCenter Server instance, you must retrieve the vCenter Server build number and deployment details from the SDDC Manager inventory. The vCenter Server instances in your system might be running different build numbers if the backups are taken during an upgrade process. You must restore each vCenter Server instance to its correct version.

Because the Management domain vCenter Server might be unavailable to authenticate the login, you use the SDDC Manager API via the shell to retrieve this information.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client.
2. Run the command to get the list of vCenter Server instances.

```
curl http://localhost/inventory/vcenters -k | json_pp
```

3. For each vCenter Server instance, record the values of these settings.

Setting	Value
domainType	Name of the domain
vmName	VM name of the vCenter Server
managementIpAddress	IP address of the vCenter Server
datastoreForVmDeploymentName	Datastore name
hostName	FQDN of the vCenter Server
version	<i>version_number-build_number</i>
Size	Size of the deployment

4. Verify that the vCenter Server version retrieved from SDDC Manager is the same as the version associated with the backup file that you plan to restore.

Retrieve the Credentials for Restoring vCenter Server

Before restoring a vCenter Server instance, you must retrieve the vCenter Server root and vCenter Single Sign-On administrator credentials from the SDDC Manager inventory. Before restoring the Management domain vCenter Server, you must also retrieve the credentials of a healthy Management domain ESXi host.

NOTE

If SDDC Manager is not operational, you can retrieve the required vCenter Server root, vCenter Single Sign-On administrator, and ESXi root credentials from the file-based backup of SDDC Manager. See [Prepare for Restoring SDDC Manager](#).

Before you can query the SDDC Manager API, you must obtain an API access token by using **admin@local** account.

1. Log in to your host machine with access to the SDDC and open a terminal.
2. Obtain the API access token.
 - a) Run the command to obtain an access token by using the **admin@local** credentials.

```
TOKEN=`curl https://<sddc_manager_fqdn>/v1/tokens -k -X POST -H "Content-Type: application/json" -d '{"username": "admin@local", "password": "admin@local_password"}' | awk -F "\"" '{print $4}'`
```

The command returns an access token and a refresh token.

- b) Record the access token.
3. Retrieve the vCenter Server **root** credentials.

- a) Run the following command to retrieve the vCenter Server **root** credentials.

```
curl https://<sddc_manager_fqdn>/v1/credentials?resourceType=VCENTER -k -X GET \
-H "Accept: application/json" -H "Authorization: Bearer $TOKEN" | json_pp
```

The command returns the vCenter Server **root** credentials.

Setting	Value
domainName	Name of the domain
resourceName	FQDN of the vCenter Server
username	root
password	<i>vcenter_server_root_password</i>

- b) Record the vCenter Server **root** credentials.

4. Retrieve the vCenter Single Sign-On administrator credentials.

- a) Run the following command to retrieve the vCenter Single Sign-On administrator credentials.

```
curl https://<sddc_manager_fqdn>/v1/credentials?resourceType=PSC -k -X GET \
-H "Accept: application/json" -H "Authorization: Bearer $TOKEN" | json_pp
```

The command returns the **administrator@vsphere.local** credentials.

Setting	Value
domainName	Name of the domain
resourceName	FQDN of the vCenter Server
username	administrator@vsphere.local
password	<i>vsphere_admin_password</i>

- b) Record the **administrator@vsphere.local** credentials.

5. If you plan to restore the management domain vCenter Server, retrieve the credentials for a healthy management domain ESXi host.

- a) Run the following command to retrieve the credentials for a management domain ESXi host.

```
curl https://<sddc_manager_fqdn>/v1/credentials?resourceType=ESXI -k -X GET \
-H "Accept: application/json" -H "Authorization: Bearer $TOKEN" | json_pp
```

The command returns the ESXi **root** credentials.

Setting	Value for first ESXi host
domainName	management domain name
resourceName	FQDN of the first ESXi host
username	root
password	<i>esxi_root_password</i>

- b) Record the ESXi **root** credentials.

Restore a vCenter Server Instance from a File-Based Backup

If a vCenter Server instance fails, you can restore it from its file-based backup. If the management domain vCenter Server and the VI workload domain vCenter Server are both in a failed state, you must restore the management domain vCenter Server before restoring the VI workload domain vCenter Server.

- Download the vCenter Server ISO file for the version of the failed instance. See [Retrieve the vCenter Server Deployment Details](#).
- If you are recovering the VI workload domain vCenter Server, verify that the management vCenter Server is available.

You deploy a new vCenter Server appliance and perform a file-based restore. If you are restoring the management domain vCenter Server, you deploy the new appliance on a healthy ESXi host in the management domain vSAN cluster. If you are restoring the VI workload domain vCenter Server, you deploy the new appliance on the management domain vCenter Server.

1. Mount the vCenter Server ISO image to your host machine with access to the SDDC and run the UI installer for your operating system.
For example, for a Windows host machine, open the `dvd-drive:\vcsa-ui-installer\win32\installer` application file.
2. Click **Restore**.
3. Complete the **Restore - Stage 1: Deploy vCenter Server** wizard.
 - a) On the **Introduction** page, click **Next**.
 - b) On the **End user license agreement** page, select the **I accept the terms of the license agreement** check box and click **Next**.
 - c) On the **Enter backup details** page, enter these values and click **Next**.

Setting	Value for vCenter Server
Location or IP/hostname	<code>sftp://sftp_server_ip/backups/vCenter/sn_vc_fqdn/backup_folder/</code>
User name	vSphere service account user
Password	<code>vsphere-service-account-password</code>

- d) On the **Review backup information** page, review the backup details, record the **vCenter Server configuration** information, and click **Next**.

You use the vCenter Server configuration information at a later step to determine the deployment size for the new vCenter Server appliance.

- e) On the **vCenter Server deployment target** page, enter the values by using the information that you retrieved during the preparation for the restore, and click **Next**.

Setting	Value for Management Domain vCenter Server	Value for VI Workload Domain vCenter Server
ESXi host or vCenter Server name	FQDN of the first ESXi host	FQDN of the management vCenter Server
HTTPS port	443	443
User name	root	administrator@vsphere.local
Password	<code>esxi_root_password</code>	<code>vsphere_admin_password</code>

- f) In the **Certificate warning** dialog box, click **Yes** to accept the host certificate.
- g) On the **Set up a target vCenter Server VM** page, enter the values by using the information that you retrieved during the preparation for the restore, and click **Next**.

Setting	Value
VM name	vCenter Server VM name
Set root password	<i>vcenter_server_root_password</i>
Confirm root password	<i>vcenter_server_root_password</i>

- h) On the **Select deployment size** page, select the deployment size that corresponds with the vCenter Server configuration information from Step 3.d and click **Next**.

Refer to vSphere documentation to map CPU count recorded from Step 3.d to a vSphere Server configuration size.

- i) On the **Select datastore** page, select these values, and click **Next**.

Setting	Value
Datastore	Datastore name
Enable thin disk mode	Selected

- j) On the **Configure network settings** page, enter the values by using the information that you retrieved during the preparation for the restore, and click **Next**.

Setting	Value
Network	Name of the vSphere distributed switch
IP version	IPv4
IP assignment	static
FQDN	FQDN of the vCenter Server
IP address	IP address of the vCenter Server
Subnet mask or prefix length	24
Default gateway	Default gateway IP address
DNS servers	DNS server IP addresses with comma separated

- k) On the **Ready to complete stage 1** page, review the restore settings and click **Finish**.

- l) When stage 1 of the restore process completes, click **Continue**.

4. Complete the **Restore - Stage 2: vCenter Server** wizard.

- a) On the **Introduction** page, click **Next**.

- b) On the **Backup details** page, in the **Encryption password** text box, enter the encryption password of the SFTP server and click **Next**.

- c) On the **Single Sign-On configuration** page, enter these values and click **Next**.

Setting	Value
Single Sign-On user name	administrator@vsphere.local
Single Sign-On password	<i>vsphere_admin_password</i>

- d) On the **Ready to complete** page, review the restore details and click **Finish**.

- e) In the **Warning** dialog box, click **OK** to confirm the restore.

- f) When stage 2 of the restore process completes, click **Close**.

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Move the Restored vCenter Server Appliance to the Correct Folder

After deploying and restoring a vCenter Server instance, you must move the new appliance to the correct folder.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > VMs and Templates**.
3. In the inventory expand **vCenter Server > Datacenter**.
4. Right-click the appliance of the restored vCenter Server instance and select **Move to folder**.
5. Select the management folder and click **OK**.

Validate the vCenter Server State

After restoring a vCenter Server instance, you must validate the state of the vCenter Server and vCenter Single Sign-On.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. In the inventory, click the management domain vCenter Server inventory, click the **Summary** tab, and verify that there are no unexpected vCenter Server alerts.
3. Click the **Linked vCenter Server systems** tab and verify that the list contains all other vCenter Server instances in the vCenter Single Sign-On domain.
4. Log in to the recovered vCenter Server instance by using a Secure Shell (SSH) client.
5. Run the command to navigate to the `bin` directory.

```
cd /usr/lib/vmware-vmware-vmware/bin
```

6. Validate the current replication status.
 - a) Run the command to list the current replication partners of the vCenter Server instance with the current replication status between the nodes.


```
./vdcadmin -f showpartnerstatus -h localhost -u administrator -w vsphere_admin_password
```
 - b) Verify that for each partner, the `vdcadmin` command output contains `Host available: Yes`, `Status available: Yes`, and `Partner is 0 changes behind`.
 - c) If you observe significant differences, because the resyncing might take some time, wait five minutes and repeat this step.
7. Repeat the procedure for the other vCenter Server instance.

Validate the SDDC Manager State After a vCenter Server Restore

After a successful vCenter Server restore, verify that the SDDC Manager inventory is consistent with the recovered VMs and that the vCenter Server instances are healthy. You use the Supportability and Serviceability tool (SoS) and the SDDC Manager patch/upgrade precheck function.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client.
2. Run the `sos` health check and verify the output.

```
sudo /opt/vmware/sddc-support/sos --health-check
```

All tests show green when SDDC Manager is in a healthy state.

3. In a Web browser, log in to SDDC Manager using the user interface.
4. In the navigation pane, click **Inventory** › **Workload Domains**.
5. For each workload domain, validate the vCenter Server status.
 - a) Click the workload domain name and click the **Updates/Patches** tab.
 - b) Click **Precheck**.
 - c) Click **View status** to review the precheck result for the vCenter Server instance and verify that the status is **Succeeded**.

Restore the Configuration of a vSphere Distributed Switch

To recover the configuration of a vSphere Distributed Switch, you can restore its settings from the configuration file that you previously exported.

This procedure restores only the vSphere Distributed Switch configuration of a vCenter Server instance. The restore operation changes the settings on the vSphere Distributed Switch back to the settings saved in the configuration file. The operation overwrites the current settings of the vSphere Distributed Switch and its port groups. The operation does not delete existing port groups that are not a part of the configuration file. The vSphere Distributed Switch configuration is part of the vCenter Server backup. If you want to restore the entire vCenter Server instance, see [Restore vCenter Server](#).

1. In a web browser, log in to the vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **Networking**.
3. In the inventory expand **vCenter Server** › **Datacenter**.
4. Expand the **Management networks** folder, right-click the distributed switch and select **Settings** › **Restore configuration**.
5. On the **Restore switch configuration** page, click **Browse**, navigate to the location of the configuration file for the distributed switch, and click **Open**.
6. Select the **Restore distributed switch and all port groups** radio-button and click **Next**.
7. On the **Ready to complete** page, review the changes and click **Finish**.
8. Repeat these steps for the other vSphere Distributed Switch.
9. Review the switch configuration to verify that it is as you expect after the restore.

Restore an NSX Manager Cluster Node

If an NSX Manager instance fails, you can restore it from its file-based backup.

- Verify that you have a valid file-based backup of the failed NSX Manager instance.
- Verify that you have the SFTP server details:
 - SFTP Server IP
 - SFTP Server Username
 - SFTP Server Password
 - Encryption Password

Prepare for Restoring an NSX Manager Cluster Node

Before restoring an NSX Manager node, you must retrieve the NSX Manager build number and deployment details, as well as the credentials from the SDDC Manager inventory.

Retrieve the NSX Manager Version from SDDC Manager

Before restoring a failed NSX Manager instance, you must retrieve its version from the SDDC Manager inventory.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. Click the domain name of the failed NSX Manager instance.

3. Click the **Update/Patches** tab.
4. Under **Current versions**, in the **NSX** panel, locate and record the **NSX upgrade coordinator** value.
5. Verify that the NSX version retrieved from SDDC Manager is the same as the version associated with the backup file that you plan to restore.

Retrieve the Credentials for Restoring NSX Manager from SDDC Manager

Before restoring a failed NSX Manager instance, you must retrieve the NSX Manager **root** and **admin** credentials from the SDDC Manager inventory.

Before you can query the SDDC Manager API, you must obtain an API access token by using an API service account.

1. Log in to your host machine with access to the SDDC and open a terminal.
2. Obtain the API access token.
 - a) Run the command to obtain an access token by using the **admin@local** account credentials.

```
curl 'https://<sddc_manager_fqdn>/v1/tokens' -k -X POST -H 'Content-Type: application/json' -H 'Accept: application/json' -d '{"username" : "service_user", "password" : "service_user_password"}'
```

The command returns an access token and a refresh token.

- b) Record the access token.
3. Retrieve the NSX Manager **root** and **admin** credentials.
 - a) Run the command to retrieve the NSX Manager **root** and **admin** credentials.

```
curl 'https://<sddc_manager_fqdn>/v1/credentials?resourceType=NSXT_MANAGER' -i -X GET -H 'Accept: application/json' -H 'Authorization: Bearer access_token'
```

The command returns the NSX Manager **root** and **admin** credentials.

- b) Record the NSX Manager **root** and **admin** credentials for the instance you are restoring.

Restore the First Node of a Failed NSX Manager Cluster

If all three NSX Manager nodes in an NSX Manager cluster are in a failed state, you begin the restore process by restoring the first cluster node.

IMPORTANT

This procedure is not applicable in use cases when there are operational NSX Manager cluster nodes.

- If two of the three NSX Manager nodes in the NSX Manager cluster are in a failed state, you begin the restore process by deactivating the cluster. See [Deactivate the NSX Manager Cluster](#).
- If only one of the three NSX Manager nodes in the NSX Manager cluster is in a failed state, you directly restore the failed node to the cluster. See [Restore an NSX Manager Node to an Existing NSX Manager Cluster](#).

Redeploy the First Node of a Failed NSX Manager Cluster

You deploy a new NSX Manager instance by using the configuration of the first NSX Manager cluster node.

- Download the NSX Manager OVA file for the version of the failed NSX Manager cluster. See [Retrieve the NSX Manager Version from SDDC Manager](#).
- Verify that the backup file that you plan to restore is associated with the version of the failed NSX Manager cluster.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > VMs and Templates**.
3. In the inventory, expand **vCenter Server > Datacenter**.
4. Right-click the NSX folder and select **Deploy OVF Template**.
5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, navigate to the location of the NSX Manager OVA file, click **Open**, and click **Next**.
6. On the **Select a name and folder** page, enter the VM name and click **Next**.
7. On the **Select a compute resource** page, click **Next**.
8. On the **Review details** page, click **Next**.
9. On the **Configuration** page, select the appropriate size and click **Next**.
For the management domain, select **Medium** and for workload domains, select **Large** unless you changed these defaults during deployment.
10. On the **Select storage** page, select the vSAN datastore, and click **Next**.
11. On the **Select networks** page, from the **Destination network** drop-down menu, select the management network distributed port group, and click **Next**.
12. On the **Customize template** page, enter these values and click **Next**.

Setting	Value for first NSX Manager cluster node
System root user password	<i>nsx_root_password</i>
CLI admin user password	<i>nsx_admin_password</i>
CLI audit user password	<i>nsx_audit_password</i>
Hostname	Enter hostname for the appliance using FQDN format.
Default IPv4 gateway	Enter the default gateway for the appliance.
Management network IPv4 address	Enter the IP Address for the appliance.
Management network netmask	Enter the subnet mask for the appliance.
DNS server list	Enter the DNS servers for the appliance.
NTP server list	Enter the NTP server for the appliance.
Enable SSH	Deselected
Allow root SSH logins	Selected

13. On the **Ready to complete** page, review the deployment details and click **Finish**.

Restore the First Node in a Failed NSX Manager Cluster from a File-Based Backup

You restore the file-based backup of the first NSX Manager cluster node to the newly deployed NSX Manager instance.

1. In a web browser, log in to the NSX Manager node for the domain by using the user interface (https://<nsx_manager_node_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left navigation pane, under **Lifecycle management**, click **Backup and restore**.
4. In the **NSX configuration** pane, under **SFTP server**, click **Edit**.
5. In the **Backup configuration** dialog box, enter these values, and click **Save**.

Setting	Value
FQDN or IP address	IP address of SFTP server
Protocol	SFTP
Port	22
Directory path	/backups
Username	Service account user name For example, svc-vcf-bck@rainpole.io
Password	<i>service_account_password</i>
SSH fingerprint	<i>SFTP_ssh_fingerprint</i>

- Under **Backup history**, select the target backup, and click **Restore**.
- During the restore, when prompted, reject adding NSX Manager nodes by clicking **I understand** and **Resume**.

A progress bar displays the status of the restore operation with the current step of the process.

Validate the Status of the First NSX Manager Cluster Node

After you restored the first NSX Manager cluster node, you validate the services state from the VM Web console of the restored node.

- In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
- Select **Menu** › **VMs and Templates**.
- In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
- Click the VM name of the newly deployed first NSX Manager cluster node, click **Launch Web Console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

- Run the command to view the cluster status.

```
get cluster status
```

The services on the single-node NSX Manager cluster appear as `UP`.

Deactivate the NSX Manager Cluster

If two of the three NSX Manager cluster nodes are in a failed state or if you restored the first node of a failed NSX Manager cluster, you must deactivate the cluster.

IMPORTANT

This procedure is not applicable in use cases when there are two operational NSX Manager cluster nodes.

If only one of the three NSX Manager nodes in the NSX Manager cluster is in a failed state, after you prepared for the restore, you directly restore the failed node to the cluster. See [Restore an NSX Manager Node to an Existing NSX Manager Cluster](#).

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **VMs and Templates**.
3. In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
4. Click the VM of the operational NSX Manager node in the cluster, click **Launch Web Console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

5. Run the command to deactivate the cluster


```
deactivate cluster
```
6. On the **Are you sure you want to remove all other nodes from this cluster? (yes/no)** prompt, enter *yes*. You deactivated the cluster.

Power off and delete the two failed NSX Manager nodes from inventory.

Restore an NSX Manager Node to an Existing NSX Manager Cluster

If only one of the three NSX Manager cluster nodes is in a failed state, you restore the failed node to the existing cluster. If two of the three NSX Manager cluster nodes are in a failed state, you repeat this process for each of the failed nodes.

Detach the Failed NSX Manager Node from the NSX Manager Cluster

Before you recover a failed NSX Manager node, you must detach the failed node from the NSX Manager cluster.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **VMs and Templates**.
3. In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
4. Click the VM of an operational NSX Manager node in the cluster, click **Launch Web Console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

5. Retrieve the UUID of the failed NSX Manager node.
 - a) Run the command to view the details of the cluster members.

```
get cluster status
```

The status of the failed node is *Down*.

- b) Record the UUID of the failed NSX Manager node.

- Run the command to detach the failed node from the cluster

```
detach node failed_node_uuid
```

The detach process might take some time.

- When the detaching process finishes, run the command to view the cluster status.

```
get cluster status
```

The status of all cluster nodes is Up.

Redeploy the Failed NSX Manager Node

You deploy a new NSX Manager instance by using the configuration of the failed node.

Download the NSX Manager OVA file for the version of the failed NSX Manager instance. See [Retrieve the NSX Manager Version from SDDC Manager](#).

- In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
- Select **Menu** > **VMs and Templates**.
- In the inventory expand **vCenter Server** > **Datacenter**.
- Right-click the NSX folder and select **Deploy OVF Template**.
- On the **Select an OVF template** page, select **Local file**, click **Upload files**, navigate to the location of the NSX Manager OVA file, click **Open**, and click **Next**.
- On the **Select a name and folder** page, in the **Virtual machine name** text box, enter VM name of the failed node, and click **Next**.
- On the **Select a compute resource** page, click **Next**.
- On the **Review details** page, click **Next**.
- On the **Configuration** page, select **Medium**, and click **Next**.
- On the **Select storage** page, select the vSAN datastore, and click **Next**.
- On the **Select networks** page, from the **Destination network** drop-down menu, select the management network distributed port group, and click **Next**.
- On the **Customize template** page, enter these values and click **Next**.

Setting	Value
System root user password	<i>nsx_root_password</i>
CLI admin user password	<i>nsx_admin_password</i>
CLI audit password	<i>nsx_audit_password</i>
Hostname	<i>failed_node_FQDN</i>
Default IPv4 gateway	Enter the default gateway for the appliance.
Management network IPv4 address	<i>failed_node_IP_address</i>
Management network netmask	Enter the subnet mask for the appliance.
DNS server list	Enter the DNS servers for the appliance.
NTP servers list	Enter the NTP services for the appliance.
Enable SSH	Deselected
Allow root SSH logins	Selected

- On the **Ready to complete** page, review the deployment details and click **Finish**.
The NSX Manager virtual machine begins to deploy.

Join the New NSX Manager Node to the NSX Manager Cluster

You join the newly deployed NSX Manager node to the cluster by using the virtual machine web console from the vSphere Client.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** > **VMs and Templates**.
3. In the inventory expand **vCenter Server** > **Datacenter** > **NSX Folder**.
4. Click the VM of an operational NSX Manager node in the cluster, click **Launch web console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

5. Retrieve the ID of the NSX Manager cluster.
 - a) Run the command to view the cluster ID.


```
get cluster config | find Id:
```
 - b) Record the cluster ID.
6. Retrieve the API thumbprint of the NSX Manager API certificate.
 - a) Run the command to view the certificate API thumbprint.


```
get certificate api thumbprint
```
 - b) Record the certificate API thumbprint.
7. Exit the VM Web console.
8. In the vSphere Client, click the VM of the newly deployed NSX Manager node, click **Launch Web console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

9. Run the command to join the new NSX Manager node to the cluster.


```
join existing_node_ip cluster-id cluster_id thumbprint api_thumbprint username admin
```

The new NSX Manager node joins the cluster.

Validate the Status of the NSX Manager Cluster

After you added the new NSX Manager node to the cluster, you must validate the operational state of the NSX Manager cluster.

To view the state of the NSX Manager cluster, you log in to the NSX Manager for the particular domain.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)

2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Appliances**.
4. Verify that the **Cluster** status is green and *Stable* and that each cluster node is *Available*.

Restore the NSX Manager SSL Certificate

After you add the new NSX Manager node to the cluster and validate the cluster status, you must restore the SSL certificate to the new node.

To view the certificate of the failed NSX Manager cluster node, you log in to the NSX Manager for the domain.

Table 197: NSX Manager Clusters in the SDDC

NSX Manager Cluster	NSX Manager URL
Management domain NSX Manager cluster	<code>https://<FQDN of management domain NSX Manager>/login.jsp?local=true</code>
Workload domain NSX Manager cluster	<code>https://<FQDN of workload domain NSX Manager>/login.jsp?local=true</code>

This procedure is an example for restoring the certificate of a management domain NSX Manager cluster node.

1. In a Web browser, log in to the NSX Manager cluster for the management domain.

Setting	Value
URL	<code>https://<FQDN of management domain NSX Manager>/login.jsp?local=true</code>
User name	<code>admin</code>
Password	<code>nsx_admin_password</code>

2. On the main navigation bar, click **System**.
3. In the left pane, under **Settings**, click **Certificates**.
4. Locate and copy the ID of the certificate that was issued by CA to the node that you are restoring.
5. Run the command to install the CA-signed certificate on the new NSX Manager node.

```
curl -H 'Accept: application/json' -H 'Content-Type: application/json' \ --insecure -u 'admin:nsx_admin_password' -X POST \ 'https://nsx_host_node/api/v1/node/services/http?action=apply_certificate&certificate_id=certificate_id'
```

IMPORTANT

If assigning the certificate fails because the certificate revocation list (CRL) verification fails, see <https://kb.vmware.com/kb/78794>. If you disable the CRL checking to assign the certificate, after assigning the certificate, you must re-enable the CRL checking.

Restart the NSX Manager Node

After assigning the certificate, you must restart the new NSX Manager node.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (`https://<vcenter_server_fqdn>/ui`).
2. Select **Menu** > **VMs and Templates**.

3. In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
4. Right click the new NSX Manager VM and select **Guest OS** › **Restart**.

Validate the Status of the NSX Manager Cluster

After restoring an NSX Manager node, you must validate the system status of the NSX Manager cluster.

To view the system status of the NSX Manager cluster, you log in to the NSX Manager for the particular domain.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the **Home** page, click **Monitoring Dashboards** › **System**.
3. Verify that all components are healthy.
4. If the host transport nodes are in a `Pending` state, run **Configure NSX** on these nodes to refresh the UI.

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Update or Recreate the VM Anti-Affinity Rule for the NSX Manager Cluster Nodes

During the NSX Manager bring-up process, SDDC Manager creates a VM anti-affinity rule to prevent the VMs of the NSX Manager cluster from running on the same ESXi host. If you redeployed all NSX Manager cluster nodes, you must recreate this rule. If you redeployed one or two nodes of the cluster, you must add the new VMs to the existing rule.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **Hosts and Clusters**.
3. In the inventory expand **vCenter Server** › **Datacenter**.
4. Click the cluster object.
5. Click the **Configure** tab and click **VM/Host Rules**.
6. Update or recreate the VM anti-affinity rule.
 - If you redeployed one or two nodes of the cluster, add the new VMs to the existing rule.
 1. Click the VM anti-affinity rule name and click **Edit**.
 2. Click **Add VM/Host rule member**, select the new NSX Manager cluster nodes, and click **Add**.
 - If you redeployed all NSX Manager cluster nodes, click **Add VM/Host rule**, enter these values to create the rule, and click **OK**.

Setting	Value
Name	Enter the name of the anti-affinity rule
Type	Separate virtual machines
Members	Click Add VM/Host rule member , select the NSX Manager cluster nodes, and click Add .

Validate the SDDC Manager Inventory State

After a successful restore of an NSX Manager cluster, you must verify that the SDDC Manager inventory is consistent with the recovered virtual machines. You run this verification by using the `sos` tool.

1. Log in to SDDC Manager by using a Secure Shell (SSH).
2. Verify the SDDC Manager health.

- a) Run the command to view the details about the VMware Cloud Foundation system.

```
sudo /opt/vmware/sddc-support/sos --get-vcf-summary
```

- b) When prompted, enter the *vcf_password*.

All tests show green state.

3. Run the command to collect the log files from the restore of the NSX Manager cluster.

```
sudo /opt/vmware/sddc-support/sos --domain-name domain_name --nsx-logs
```

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Restoring NSX Edge Cluster Nodes

If one or both NSX Edge cluster nodes fail due to a hardware or software issue, you must redeploy the failed NSX Edge instances. You do not restore the NSX Edge nodes from a backup.

Prepare for Restoring NSX Edge Cluster Nodes

Before restoring an NSX Edge node, you must retrieve its deployment details from the NSX Manager cluster and retrieve the credentials of the failed NSX Edge node from SDDC Manager.

Retrieve the NSX Edge Node Deployment Details from NSX Manager Cluster

Before restoring a failed NSX Edge node, you must retrieve its deployment details from the NSX Manager cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric** > **Nodes**.
4. Click the **Edge Transport Nodes** tab.
5. Select the check-box for the failed NSX Edge node.
6. Click **Actions** and select **Change node settings**.
7. Record the **Host name/FQDN** value and click **Cancel**.
8. Click **Actions** and select **Change Edge VM Resource Reservations**.
9. Record the **Existing form factor** value and click **Cancel**.
10. Click the name of the NSX Edge node that you plan to replace and record the following values.
 - Name
 - Management IP
 - Transport Zones
 - Edge Cluster
11. Click **Edit**, record the following values, and click **Cancel**.
 - Edge Switch Name
 - Uplink Profile
 - IP Assignment
 - Teaming Policy Uplink Mapping

Retrieve the NSX Edge Node Credentials from SDDC Manager

Before restoring the failed NSX Edge node that is deployed by SDDC Manager, you must retrieve its credentials from the SDDC Manager inventory.

1. In the SDDC Manager user interface, from the navigation pane click **Developer center**.
2. Click the **API explorer** tab.
3. Expand **APIs for managing credentials** and click **GET /v1/credentials**.
4. In the **resourceName** text box, enter the FQDN of the failed NSX Edge node, and click **Execute**.
5. Under **Response**, click **PageOfCredential** and click each credential ID.
6. Record the user names and passwords for these credentials.

Credential Type	Username	Password
SSH	root	<i>edge_root_password</i>
API	admin	<i>edge_admin_password</i>
AUDIT	audit	<i>edge_audit_password</i>

Retrieve the Workload Domain vSphere Cluster ID from SDDC Manager

If you are restoring a failed workload domain NSX Edge node, you must retrieve the ID of the vSphere cluster for the workload domain. During the restore process, you use this vSphere cluster ID to recreate the vSphere DRS rule name with its original name.

You use the SDDC Manager user interface to retrieve the ID of the vSphere cluster for the workload domain.

1. In the SDDC Manager user interface, from the navigation pane click **Developer center**.
2. Click the **API explorer** tab.
3. Expand **APIs for managing clusters**, click **GET /v1/clusters**, and click **Execute**.
4. Under **Response**, click **PageOfClusters** and click **Cluster**.
5. Record the **ID of the cluster** for the workload domain cluster ID.

Replace the Failed NSX Edge Node with a Temporary NSX Edge Node

You deploy a temporary NSX Edge node in the domain, add it to the NSX Edge cluster, and then delete the failed NSX Edge node.

Deploy a Temporary NSX Edge Node

To avoid conflicts with the failed NSX Edge node, you deploy a temporary NSX Edge node with a new FQDN and IP address.

Allocate the FQDN and IP address for the temporary NSX Edge node for the domain of the failed node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Click **Add edge VM**.
6. On the **Name and description** page, enter these values and click **Next**.

Setting	Value
Name	Enter the VM name
Host name/FQDN	Enter the FQDN
Form factor	Medium

7. On the **Credentials** page, enter these values and the passwords recorded in the earlier steps and then click **Next**.

Setting	Value
CLI user name	admin
CLI password	<i>edge_admin_password</i>
CLI confirm password	<i>edge_admin_password</i>
Allow SSH login	Yes
System root password	<i>edge_root_password</i>
System root password confirm	<i>edge_root_password</i>
Allow root SSH login	No
Audit user name	audit
Audit password	<i>edge_audit_password</i>
Audit confirm password	<i>edge_audit_password</i>

8. On the **Configure deployment** page, select the following and click **Next**.

Setting	Value
Compute manager	Enter the vCenter Server FQDN
Cluster	Select the cluster
Datastore	Select the vSAN datastore

9. On the **Configure node settings** page, enter these values and click **Next**.

Setting	Value
IP Assignment	Static
Management IP	Enter the management IP address.
Default Gateway	Enter the default gateway
Management interface	Select the management network distributed port group
Search domain names	Enter the search domain
DNS servers	Enter the DNS servers
NTP Servers	Enter the NTP servers

10. On the **Configure NSX** page, enter these values which are already recorded and click **Finish**.

Setting	Value
Edge switch name	Enter the edge switch name.
Transport zone	Enter the transport zone names.
Uplink profile	Enter the uplink profile name.
IP assignment	Use static IP list
Static IP list	Enter the static IP list.
Gateway	Enter the gateway IP
Subnet mask	Enter the subnet mask
Teaming policy switch mapping	Enter the values for Uplink1 and Uplink2.

Replace the Failed NSX Edge Node with the Temporary NSX Edge Node

You add the temporary NSX Edge node to the NSX Edge cluster by replacing the failed NSX Edge node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge clusters** tab.
5. Select the check-box for the NSX Edge cluster.
6. Click **Action** and select **Replace edge cluster member**.
7. From the **Replace** drop down menu, select the Failed edge node and from the **with** drop down menu, select the Temporary edge node and then click **Save**.

Delete the Failed NSX Edge Node from the NSX Manager Cluster

After replacing the failed NSX Edge node with the temporary NSX Edge node in the NSX Edge cluster, you delete the failed node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Select the check-box for the failed NSX Edge node and click **Delete**.
6. In the confirmation dialog box, click **Delete**.

Validate the Temporary State of the NSX Edge Cluster Nodes

After replacing the failed NSX Edge node with a temporary NSX Edge node, you must verify the state of the NSX Edge cluster nodes.

You validate the state of the temporary NSX Edge node and the second NSX Edge node in the cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.

5. Verify all edge transport nodes show these values.

Setting	Value
Configuration state	Success
Node status	Up
Tunnels	Upward arrow mark with number of tunnels

Replace the Temporary NSX Edge Node with the Redeployed NSX Edge Node

After you replaced and deleted the failed NSX Edge node, to return the NSX Edge cluster to its original state, you redeploy the failed node, add it to the NSX Edge cluster, and delete then temporary NSX Edge node.

Redeploy the Failed NSX Edge Node

You deploy a new NSX Edge node by using the configurations of the failed NSX Edge node that you retrieved during the preparation for the restore.

To return the NSX Edge cluster to the original state, you must use the FQDN and IP address of the failed NSX Edge node that you deleted. This procedure ensures that the inventory in SDDC Manager is accurate.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Click **Add edge VM**.
6. On the **Name and description** page, enter these values and click **Next**.

Setting	Value
Name	Enter the VM name
Host name/FQDN	Enter the FQDN
Form factor	Medium

7. On the **Credentials** page, enter these values which are recorded earlier and click **Next**.

Setting	Value
CLI user name	admin
CLI password	<i>edge_admin_password</i>
CLI confirm password	<i>edge_admin_password</i>
Allow SSH login	Yes
System root password	<i>edge_root_password</i>
System root password confirm	<i>edge_root_password</i>
Allow root SSH login	No
Audit user name	audit
Audit password	<i>edge_audit_password</i>

Table continued on next page

Continued from previous page

Setting	Value
Audit confirm password	<code>edge_audit_password</code>

8. On the **Configure deployment** page, select these values and click **Next**.

Setting	Value
Compute manager	Enter the vCenter Server FQDN
Cluster	Enter the cluster name
Resource pool	Enter the resource pool
Datastore	Enter the datastore

9. On the **Configure Node Settings** page, enter these values and click **Next**.

Setting	Value
IP assignment	Static
Management IP	Enter the management IP address.
Default gateway	Enter the default gateway
Management interface	Select the management network distributed port group
Search domain names	Enter the search domain
DNS servers	Enter the DNS servers
NTP servers	Enter the NTP servers

10. On the **Configure NSX** page, enter these values which are recorded earlier and click **Finish**.

Setting	Value
Edge switch name	Enter the edge switch name.
Transport zone	Enter the transport zone names.
Uplink profile	Enter the uplink profile name.
IP assignment	Use static IP list
Static IP list	Enter the static IP list.
Gateway	Enter the gateway IP
Subnet mask	Enter the subnet mask
Teaming policy switch mapping	Enter the values for Uplink1 and Uplink2.

Replace the Temporary NSX Edge Node with the Redeployed NSX Edge Node

After deploying the new NSX Edge node with the same configuration as the failed NSX Edge node, you replace the temporary NSX Edge node with the redeployed failed node in the NSX- Edge cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.

4. Click the **Edge clusters** tab.
5. Select the check-box for the NSX Edge cluster.
6. Click **Action** and select **Replace edge cluster member**.
7. From the **Replace** drop down menu, select the temporary node and from the **with** drop down menu, select the new node and then click **Save**.

Delete the Temporary NSX Edge Node

After replacing the temporary NSX Edge node with the new NSX Edge node in the NSX Edge cluster, you delete the temporary node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes > .**
4. Click the **Edge transport nodes** tab.
5. Select the check-box for the temporary NSX Edge node and click **Delete**.
6. In the confirmation dialog box, click **Delete**.

Update or Recreate the VM Anti-Affinity Rule for the NSX Edge Cluster Nodes

During the NSX Edge deployment process, SDDC Manager creates a VM anti-affinity rule to prevent the nodes of the NSX Edge cluster from running on the same ESXi host. If you redeployed the two NSX Edge cluster nodes, you must recreate this rule. If you redeployed one node of the cluster, you must add the new VM to the existing rule.

1. In a web browser, log in to the domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > Hosts and Clusters**.
3. In the inventory expand **vCenter Server > Datacenter**.
4. Click the cluster object.
5. Click the **Configure** tab and click **VM/Host Rules**.
6. Update or recreate the VM anti-affinity rule.
 - If you redeployed one of the nodes in the NSX Edge cluster, add the new VM to the existing rule.
 1. Click the VM anti-affinity rule name and click **Edit**.
 2. Click **Add VM/Host rule member**, select the new NSX Edge cluster node, and click **Add**.
 - If you redeployed the two nodes in the NSX Edge cluster, click **Add VM/Host rule**, enter these values to create the rule, and click **OK**.

Setting	Value
Name	Enter the name of the anti-affinity rule
Type	Separate virtual machines
Members	Click Add VM/Host rule member , select the NSX Edge cluster nodes, and click Add .

Validate the State of the NSX Edge Cluster Nodes

After replacing the temporary NSX Edge node with the redeployed failed NSX Edge node, you must verify the state of the NSX Edge cluster nodes.

You validate the state of the redeployed NSX Edge node and the second NSX Edge node in the cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Verify all edge transport nodes show these values.

Setting	Value
Configuration state	Success
Node status	Up
Tunnels	Upward arrow mark with number of tunnels

Image-Based Backup and Restore of VMware Cloud Foundation

For an image-based backup of the VMware Cloud Foundation, use a solution compatible with the VMware vSphere Storage APIs - Data Protection (formerly known as VMware vStorage APIs for Data Protection or VADP).

vSphere Storage APIs - Data Protection compatible backup software connects to the vCenter servers in the management domain to perform backups. In the event of failure, the backup software connects to the vCenter servers in the management domain to restore the VMs. If the management domain is lost, the vCenter servers are no longer available and must be restored first. Choosing a backup software that supports Direct Restore to an ESXi host allows restoring the vCenter Servers.

Connect your backup solution with the management domain vCenter Server and configure it. To reduce the backup time and storage cost, use incremental backups in addition to the full ones.

Acquiesced backups are enabled for VMware Aria Suite Lifecycle and Workspace ONE Access.

VMware Cloud Foundation Glossary

In VMware Cloud Foundation, you perform specific operations and use unique constructs for automated SDDC deployment and maintenance.

Term	Description
availability zone	A collection of infrastructure components. Each availability zone is isolated from the other availability zones to prevent the propagation of failure or outage across the data center. In VMware Cloud Foundation, you implement availability of workloads across availability zones by using vSAN stretched clusters.
Application virtual networks (AVNs)	Virtual networks backed by overlay or VLAN NSX segments using the encapsulation protocol of VMware NSX. An AVN uses a single IP address space to span across data centers.
bring-up	Deployment and initial configuration of a VMware Cloud Foundation system. During the bring-up process, the management domain is created and the VMware Cloud Foundation software stack is deployed on the management domain.
commission a host	Adding a host to VMware Cloud Foundation inventory. The host becomes unassigned.
dirty host	A host that has been removed from a cluster in a workload domain. A dirty host cannot be assigned to another workload domain until it is decommissioned, re-imaged, and commissioned again.

Table continued on next page

Continued from previous page

Term	Description
decommission a host	Removing an unassigned host from the VMware Cloud Foundation inventory. SDDC Manager does not manage decommissioned hosts.
NSX Edge cluster	A logical grouping of NSX Edge nodes. These nodes run on a vSphere cluster, and provide north-south and east-west routing and network services for the management or VI workload domain.
free pool	Hosts in the VMware Cloud Foundation inventory that are not assigned to a workload domain.
host	A server that is imaged with the ESXi software.
install bundle	Contains software for VI workload domains and VMware Aria Suite Lifecycle. You can use an install bundle to deploy later versions of the software components in a new VI workload domain than the versions in the Bill of Materials for VMware Cloud Foundation.
inventory	Logical and physical entities managed by VMware Cloud Foundation.
Kubernetes - Workload Management	With Kubernetes - Workload Management, you can deploy and operate the compute, networking, and storage infrastructure for vSphere IaaS Control Plane workloads. A vSphere IaaS Control Plane workload is an application with containers running inside vSphere pods, regular VMs, or Tanzu Kubernetes clusters.
Lifecycle Manager (LCM)	Automates patching and upgrading of the software stack.
management domain	One or more vSphere clusters of physical hosts that contain the management component VMs, such as vCenter Server, NSX Manager cluster, management NSX Edge cluster, SDDC Manager, and so on. The management domain supports only vSAN storage.
network pool	Automatically assigns static IP addresses to vSAN and vMotion VMkernel ports so that you don't need to enter IP addresses manually when creating a VI workload domain or adding a host or cluster to a workload domain.
update bundle	Contains software to update the VMware Cloud Foundation components in your management or VI workload domain.
principal storage	Required for each vSphere cluster, containing the data of the virtual machines in the cluster. For the management domain, only vSAN principal storage is supported. For a VI workload domain, you set the principal storage when creating the domain or when adding a cluster to the domain. You cannot change the principal storage later. See also <i>supplemental storage</i> .
SDDC Manager	A software component that provisions, manages, and monitors the logical and physical resources of a VMware Cloud Foundation system. SDDC Manager provides the user interface for managing VMware Cloud Foundation, CLI-based administrator tools, and an API for further automation.
server	A bare-metal server in a physical rack. After imaging, it is referred to as a host.
supplemental storage	Extends the capacity of the workload domain for hosting more virtual machines or storing supporting data, such as backups. You can add or remove supplemental storage to clusters in the management or VI workload domain at any time.
unassigned host	A host in the free pool that does not belong to a workload domain.
vSphere Lifecycle Manager (vLCM)	A vCenter Server service, which is integrated with VMware Cloud Foundation, that enables centralized and simplified life cycle management of ESXi hosts.
virtual infrastructure (VI) workload domain	One or more vSphere clusters that contain customer workloads. VMware Cloud Foundation scales and manages the life cycle of each VI workload domain independently. The vCenter Server instance and NSX Manager cluster for a VI workload

Table continued on next page

Continued from previous page

Term	Description
	domain are physically located in the management domain, while the NSX edge nodes - on the VI workload domain.
vSphere Lifecycle Manager baseline	A grouping of multiple bulletins. You can attach a baseline to an ESXi host and check the compliance of the host against the associated baseline. According to the type of content, baselines are patch baselines, extension baselines, and upgrade baselines. SDDC Manager creates the required baseline and baseline group for updating a cluster in a workload domain.
vSphere Lifecycle Manager image	A precise description of the software, components, vendor add-ons, and firmware to run on an ESXi host. You set up a single image and apply it to all hosts in a cluster, thus ensuring cluster-wide host image homogeneity.
workload domain	<p>A policy-based resource container with specific availability and performance attributes that combines vSphere, storage (vSAN, NFS, VMFS on FC, or vVols) and networking (VMware NSX) into a single consumable entity. A workload domain can be created, expanded, and deleted as part of the SDDC life cycle operations. It can contain clusters of physical hosts with a corresponding vCenter Server instance to manage them.</p> <p>VMware Cloud Foundation supports two types of workload domains - the management domain and one or more VI workload domains.</p>

Operations Guide

Best practices and step-by-step instructions about operating VMware Cloud Foundation™ (also called VCF) including full-stack shutdown and startup and verifying whether the state of VMware Cloud Foundation is intact after a maintenance operation.

This guide covers all software products and workload domain types that are supported by VMware Cloud Foundation including VMware vSphere® with VMware Tanzu® and VMware Aria Suite Lifecycle™.

You can follow industry best practices when performing operations in a VMware Cloud Foundation deployment. See [Best Practices for Operating VMware Cloud Foundation](#).

To maintain component integration and avoid operation faults, you follow the specified order and steps to shut down and then start up the management components in VMware Cloud Foundation. See [Shutdown and Startup of VMware Cloud Foundation](#).

To meet the requirements of your organization for security and compliance for your VMware Cloud Foundation environment including industry compliance standards, you configure manually the password policies of the individual management components in the environment. See [Password Policy Configuration for VMware Cloud Foundation](#).

Because SDDC Manager does not manage certificates for ESXi hosts, if required by the policy of your organization, you manually replace the default host certificates that are signed by the VMware Certificate Authority (VMCA) with external CA-signed certificates. See [ESXi Certificate Management for VMware Cloud Foundation](#).

Instead of the default step-by-step approach by using product user interface, you can manage the certificates of the management components for the management domain or a VI workload domain in an automated way by running a PowerShell script. See [Managing Management Component Certificates in VMware Cloud Foundation by Using PowerShell](#).

Intended Audience

The information in *VMware Cloud Foundation Operations Guide* is intended for data center cloud administrators and operators who are familiar with:

- Concepts of virtualization and software-defined data centers (SDDCs)
- Networking and concepts such as uplinks, NICs, and IP networks
- Hardware components such as top-of-rack (ToR) switches, inter-rack switches, servers with direct attached storage, cables, and power supplies
- Methods for setting up physical racks in a data center
- Using VMware vSphere® to work with virtual machines.

PowerShell Modules for VMware Cloud Foundation Operations

As an alternative to step-by-step workflows in the product UI, you can perform VMware Cloud Foundation operations in an automated infrastructure-as-code approach by using open-source PowerShell modules from VMware.

Table 198: PowerShell Modules for VMware Cloud Foundation Operations

Operation Type	PowerShell Module or Script	More Information
Shutdown and startup of workload domains	VMware.CloudFoundation.PowerManagement	See VMware.CloudFoundation.PowerManagement open-source project in GitHub .

Table continued on next page

Continued from previous page

Operation Type	PowerShell Module or Script	More Information
Password policy configuration of management components	VMware.CloudFoundation.PasswordManagement	See VMware.CloudFoundation.PasswordManagement module in PowerShell Gallery .
Additional certificate management	VMware.CloudFoundation.CertificateManagement	See VMware.CloudFoundation.CertificateManagement open-source project in Github .
PowerShell based interaction with the VMware Cloud Foundation API	PowerVCF	PowerVCF open-source project in Github
Report-based health monitoring of VMware Cloud Foundation	VMware.CloudFoundation.Reporting	VMware.CloudFoundation.Reporting open-source project in Github
Automation for VMware Validated Solutions	PowerValidatedSolutions	PowerValidatedSolutions open-source project in Github
Automation for SDDC Manager bundle management	PowerShell Script for VMware Cloud Foundation Bundle Management	Knowledge Base article 94760

Related VMware Cloud Foundation Publications

The [VMware Cloud Foundation 5.2 Release Notes](#) lists the software components, new features, compatibility, and known issues in VMware Cloud Foundation.

- The [Planning and Preparation Workbook](#) contains the environment specification of your VMware Cloud Foundation deployment. It also provides dynamic sizing guidance.
- The [Design Guide](#) explains the design principles of and provides best practices for the management component configuration in a VMware Cloud Foundation environment.
- The [Deployment Guide](#) is intended for data center cloud administrators who deploy a VMware Cloud Foundation system in their organization's data center.
- The [Administration Guide](#) contains detailed information about how to administer and operate a VMware Cloud Foundation system in your data center.
- The [Lifecycle Management Guide](#) document describes how to manage the life cycle of a *VMware Cloud Foundation* environment.
- *VMware Validated Solutions* provide technical reference for designing and implementing add-on configurations on top of VMware Cloud Foundation that solve a business use case, such as, central identity management, workload provisioning, vSphere with Tanzu configuration, and others.

VMware Cloud Foundation Glossary

The [VMware Cloud Foundation Glossary](#) defines terms specific to VMware Cloud Foundation.

Shutdown and Startup of VMware Cloud Foundation

Shutting down VMware Cloud Foundation, for example, during hardware or power maintenance, and then starting it up must be done in a way that prevents data loss or appliance malfunction, and supports collection of troubleshooting data.

You follow a strict order and steps for shutdown and startup of the VMware Cloud Foundation management components.

Shutting Down and Starting Up VMware Cloud Foundation by Using PowerShell

Instead of the default step-by-step approach by using product user interface, you can shut down the management domain or a VI workload domain in an automated way by running a PowerShell script.

To shut down or start up the management domain or a VI workload domain, you run sample PowerShell scripts that come with the [VMware.CloudFoundation.CertificateManagement module in PowerShell Gallery](#). The scripts follow the order for manual shutdown and startup of VMware Cloud Foundation. You can complete the workflow manually at any point. You can also run the scripts multiple times.

To read the documentation, provide feedback, report an issue with automation, or contribute to the `VMware.CloudFoundation.PowerManagement` module, go to the [VMware.CloudFoundation.PowerManagement open-source project in GitHub](#).

Shutting Down VMware Cloud Foundation

To avoid data loss and maintain the SDDC components operational, you follow a specific order when shutting down the management virtual machines in VMware Cloud Foundation.

- Verify that you have complete backups of all management components.
- Verify that the management virtual machines are not running on snapshots.
- If a vSphere Storage APIs for Data Protection (VADP) based backup solution is running on the management clusters, verify that the solution is properly shut down by following the vendor guidance.
- To reduce the startup time before you shut down the management virtual machines, migrate the VMware vCenter Server® instance for the management domain to the first VMware ESXi™ host in the default management cluster in the management domain.

You shut down the customer workloads and the management components for the VI workload domains before you shut down the components for the management domain.

If the VMware NSX® Manager™ cluster and VMware NSX® Edge™ cluster are shared with other VI workload domains, shut down the NSX Manager and NSX Edge clusters as part of the shutdown of the first VI workload domain.

Starting Up VMware Cloud Foundation

To maintain the components integration and avoid operation faults, you follow a specified order to start up the management virtual machines in VMware Cloud Foundation.

- Verify that external services such as Active Directory, DNS, NTP, SMTP, and FTP or SFTP are available.
- If a vSphere Storage APIs for Data Protection (VADP) based backup solution is deployed on the default management cluster, verify that the solution is properly started and operational according to the vendor guidance.

You start the management components for the management domain first. Then, you start the management components for the VI workload domains and the customer workloads.

If the NSX Manager cluster and NSX Edge cluster are shared with other VI workload domains, start the other VI workload domains first. Start up NSX Manager and NSX Edge nodes as part of the startup of the last workload domain.

Password Policy Configuration for VMware Cloud Foundation

Configuring password policies includes the configuration of password expiration, complexity and account lockout policies according to the requirements of your organization which might be based on industry compliance standards. In VMware Cloud Foundation, this activity is performed manually.

Password Policy Configuration and Password Management

VMware Cloud Foundation does not prescribe or automate the process of configuring a password policy across the system. However, your organization might have specific requirements defined either by the organization itself or through an industry compliance standard that prescribes the changes that you must make to the default policy configuration.

After you configure the password policy, you can use SDDC Manager to rotate or manually update the passwords of the management components in VMware Cloud Foundation by using automation. See [Managing Passwords in VMware Cloud Foundation](#) in *VMware Cloud Foundation Administration Guide*.

For information about password policy design including the details and justification for the configuration of password expiration, complexity and account lockout policies, see *Information Security and Access Control Design* in the *Identity and Access Management for VMware Cloud Foundation* validated solution.

Table 199: Password Policies Support in the Management Components of VMware Cloud Foundation

Password Policy	Support by Management Component
Password expiration	<ul style="list-style-type: none"> • ESXi • vCenter Single Sign-On • vCenter Server • NSX Manager • NSX Edge • SDDC Manager
Password complexity	<ul style="list-style-type: none"> • ESXi • vCenter Single Sign-On • vCenter Server • NSX Manager • NSX Edge • SDDC Manager
Account lockout	<ul style="list-style-type: none"> • ESXi • vCenter Single Sign-On • vCenter Server • NSX Manager • NSX Edge • SDDC Manager

Manual and Automated Password Policy Configuration

To configure password policies in VMware Cloud Foundation, you can follow a step-by-step approach by using product user interface or an automated approach by running PowerShell commands that are available in the [VMware.CloudFoundation.PowerManagement module in PowerShell Gallery](#).

If you want to learn more details about, provide feedback, report an issue with automation, or contribute to the `VMware.CloudFoundation.PasswordManagement` module, go to the [VMware.CloudFoundation.PasswordManagement open-source project in GitHub](#).

Approaches to Password Policy Configuration

For initial configuration of the password policy in VMware Cloud Foundation, you usually configure all password policies on a management component and then proceed with the next one. You can also configure a specific property in a password policy across several management components.

Table 200: Password Policy Configuration by Management Component

Management Component	
ESXi	<ul style="list-style-type: none"> • Configure the Local User Password Expiration Policy for ESXi • Configure the Local User Password Complexity Policy for ESXi

Table continued on next page

Continued from previous page

Management Component	
	<ul style="list-style-type: none"> • Configure the Local Account Lockout Policy for ESXi
vCenter Single Sign-on	<ul style="list-style-type: none"> • Configure the Password Expiration Policy for vCenter Single Sign-On • Configure the Password Complexity Policy for vCenter Single Sign-On • Configure the Account Lockout Policy for vCenter Single Sign-On
vCenter Server	<ul style="list-style-type: none"> • Password expiration policy <ul style="list-style-type: none"> – Configure the Global Password Expiration Policy for vCenter Server – Configure the root User Password Expiration Policy for vCenter Server • Configure the Local User Password Complexity Policy for vCenter Server • Configure the root User Account Lockout Policy for vCenter Server
NSX Manager	<ul style="list-style-type: none"> • Configure the Local User Password Expiration Policy for NSX Manager • Configure the Local User Password Complexity Policy for NSX Manager • Configure the Local User Account Lockout Policy for NSX Manager
NSX Edge	<ul style="list-style-type: none"> • Configure the Local User Password Expiration Policy for NSX Edge • Configure the Local User Password Complexity Policy for NSX Edge • Configure the Local User Account Lockout Policy for NSX Edge
SDDC Manager	<ul style="list-style-type: none"> • Configure the Local User Password Expiration Policy for SDDC Manager • Configure the Local User Password Complexity Policy for SDDC Manager • Configure the Local User Account Lockout Policy for SDDC Manager

Prerequisites

To perform the configuration associated with password policy configuration, verify that your system fulfills the following prerequisites.

Category	Prerequisite
Environment	<ul style="list-style-type: none"> • Verify that your VMware Cloud Foundation instance is healthy and fully operational.
Infrastructure-as-code	To use the infrastructure-as-code method for password policy configuration, verify that your system fulfills the prerequisites, described in the documentation of the

Table continued on next page

Continued from previous page

Category	Prerequisite
	VMware.CloudFoundation.PasswordManagement open-source project in GitHub.

Lifecycle Management Guide

Manage the lifecycle of a VMware Cloud Foundation environment, including downloading installation and upgrade bundles, upgrading the management domain and VI workload domains, and monitoring your upgrade.

Upgrading VMware Cloud Foundation to 5.2.x

This *VMware Cloud Foundation Lifecycle Management* document describes how to manage the lifecycle of a VMware Cloud Foundation environment. The information includes prerequisites, step-by-step configuration instructions, and suggested best practices.

NOTE

Review the [VMware Interoperability Matrix](#) to verify compatibility and upgradability before planning and starting an upgrade.

You can perform a sequential or skip-level upgrade to VMware Cloud Foundation 5.2.x from VMware Cloud Foundation 4.5 or later. If your environment is at a version earlier than 4.5, you must upgrade the management domain and all VI workload domains to VMware Cloud Foundation 4.5 or later before you can upgrade to VMware Cloud Foundation 5.2.x.

CAUTION

vSphere with Tanzu enabled clusters, may require a specific upgrade sequence. See [KB 92227](#) for more information.

The first step is to download the bundles for each VMware Cloud Foundation component that requires an upgrade. After all of the bundles are available in SDDC Manager, upgrade the management domain and then your VI workload domains.

- [Upgrade the Management Domain to VMware Cloud Foundation 5.2.x](#)
- [Upgrade VI Workload Domains to VMware Cloud Foundation 5.2.x](#)

SDDC Manager Functionality During an Upgrade to VMware Cloud Foundation 5.2

During the upgrade to VMware Cloud Foundation 5.2, some SDDC Manager functionality may be limited during each phase of the upgrade. Prior to initiating the upgrade determine if you will need to perform any of these tasks.

Upgrade States and Terminology

- **Source BOM** - Prior to initiating the upgrade all components are at VMware Cloud Foundation 4.5.x, 5.0, or 5.1.
- **SDDC Manager only** - You have updated SDDC Manager to 5.2, but none of the other BOM components.
- **Split BOM** - Management domain or VI Workload Domain is only partially updated to VMware Cloud Foundation 5.2.
- **Mixed 4.5.x/5.x BOM** - Some workload domains (Management or VI) have been completely upgraded to VMware Cloud Foundation 5.2 and at least one VI Workload Domain is at the Source 4.5.x BOM version.
- **Mixed 5.x BOM** - Some workload domains (Management or VI) have been completely upgraded to VMware Cloud Foundation 5.2 and at least one VI Workload Domain is at the Source 5.0 or 5.1 BOM version.
- **Target BOM** - All components are at VMware Cloud Foundation 5.2.

When a VMware Cloud Foundation instance is in Source BOM or Target BOM, the features available within SDDC Manager are as expected for that given release. However when in a Mixed BOM the operations available vary per workload domain depending on which state the domain itself is in.

The following table indicates the functions available within SDDC Manager during an upgrade.

Table 201: SDDC Manager Functionality During Upgrade

Category	Feature	SDDC Manager only	Split BOM	Mixed 4.5.x/5.x BOM	Mixed 5.x BOM
Backup / Restore	Configure and perform Backup / Restore	Y	Y	Y	Y
CEIP	Activate / Deactivate CEIP	Y	Y	Y	Y
Certificate Management	View/Generate/ Upload/Install	Y	Y	Y	Y
NSX Edge Cluster	Expand edge cluster	Y	Y	Y	Y
DNS / NTP configuration	Validate / Configure DNS	Y	Y	Y	Y
	Validate / Configure NTP	Y	Y	Y	Y
Hosts	Commission / Decommission Host	Y	Y	Y	Y
Licensing	Update License Key Information	Y	Y	Y	Y
	Add License Key	Y	Y	Y	Y
	Relicensing	Y	Y	Y	Y
	License check	Y	Y	Y	Y
LCM	Connect to VMware or Dell Depot / Download Bundles	Y	Y	Y	Y
	LCM Pre checks	Y	Y	Y	Y
	Schedule Bundle Download	Y	Y	Y	Y
	Install vCenter Patch	Y	Y	Y	Y
	Install ESXi Patch	Y	Y	Y	Y
	Install NSX Patch	Y	Y	Y	Y
Networking	Create / Edit / Delete Network Pool	Y	Y	Y	Y
Password Management	Rotate/Update/ Retry/Cancel	Y	Y	Y	Y
User Operations	Add / Remove User / Group	Y	Y	Y	Y
Workload Domain	Add/Remove ESXi Host	Y	Y	Y	Y
	Add/Remove vSphere Cluster	Y	Y	Y	Y

Table continued on next page

Continued from previous page

Category	Feature	SDDC Manager only	Split BOM	Mixed 4.5.x/5.x BOM	Mixed 5.x BOM
	Add 4.5.x Workload Domain	Y If the management domain is at 4.5.x. NOTE Contact Broadcom Support for a workload and if the management domain is at 5.x.	Y If the management domain is at 4.5.x. NOTE Contact Broadcom Support for a workload and if the management domain is at 5.x.	Y If the management domain is at 4.5.x. NOTE Contact Broadcom Support for a workload and if the management domain is at 5.x.	N/A
	Add 5.x Workload Domain in ELM mode	Y	Y	Y	Y
	Add 5.x Isolated Workload Domain	Y	Y	Y	Y
	Remove 4.5.x Workload Domain	Y	Y	Y	N/A
	Remove 5.0 Workload Domain	Y	Y	Y	Y
	Stretch a vSphere Cluster	You cannot stretch clusters in 4.5.x workload domains, but can stretch cluster in 5.x workload domains.	You cannot stretch clusters in 4.5.x workload domains, but can stretch cluster in 5.x workload domains.	You cannot stretch clusters in 4.5.x workload domains, but can stretch cluster in 5.x workload domains.	Y
	Unstretch a vSphere Cluster	You cannot unstretch clusters in 4.5.x workload domains, but can unstretch cluster in 5.x workload domains.	You cannot unstretch clusters in 4.5.x workload domains, but can unstretch cluster in 5.x workload domains.	You cannot unstretch clusters in 4.5.x workload domains, but can unstretch cluster in 5.x workload domains.	Y
	Expand a Stretched vSphere Cluster	You cannot expand clusters in 4.5.x workload domains, but can expand clusters in 5.x workload domains.	You cannot expand clusters in 4.5.x workload domains, but can expand clusters in 5.x workload domains.	You cannot expand clusters in 4.5.x workload domains, but can expand clusters in 5.x workload domains.	Y

Table continued on next page

Continued from previous page

Category	Feature	SDDC Manager only	Split BOM	Mixed 4.5.x/5.x BOM	Mixed 5.x BOM
	Shrink a Stretched vSphere Cluster	You cannot shrink clusters in 4.5.x workload domains, but can shrink clusters in 5.x workload domains.	You cannot shrink clusters in 4.5.x workload domains, but can shrink clusters in 5.x workload domains.	You cannot shrink clusters in 4.5.x workload domains, but can shrink clusters in 5.x workload domains.	Y

vSphere UI Client Plug-ins

Identify all vSphere UI client plug-ins prior to the upgrade.

It may be possible to upgrade some vSphere UI client plug-ins before upgrading to vSphere 8.0. Contact your 3rd Party vendor to determine the best upgrade path.

Monitor VMware Cloud Foundation Updates

You can monitor in-progress updates for VMware Cloud Foundation components.

1. In the In-Progress Updates section, click **View Status** to view the high-level update progress and the number of components to be updated.
2. Details of the component being updated is shown below that. The image below is an example and may not reflect the actual versions.

VMware Cloud Foundation Update Status

VMware Cloud Foundation Update 4.0.1.0

Released 06/23/2020 11 GB

This VMware Cloud Foundation Upgrade 4.0.0.1 to 4.0.1.0 contains features, critical bugs and security fixes
<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.01/rn/VMware-Cloud-Foundation-401-Release-N>

Updating COMMON SERVICES

> SDDC MANAGER

 In Progress

- Click the arrow to see a list of tasks being performed to update the component. As the task is completed, it shows a green check mark.

VMware Cloud Foundation Update Status

VMware Cloud Foundation Update 4.0.1.0

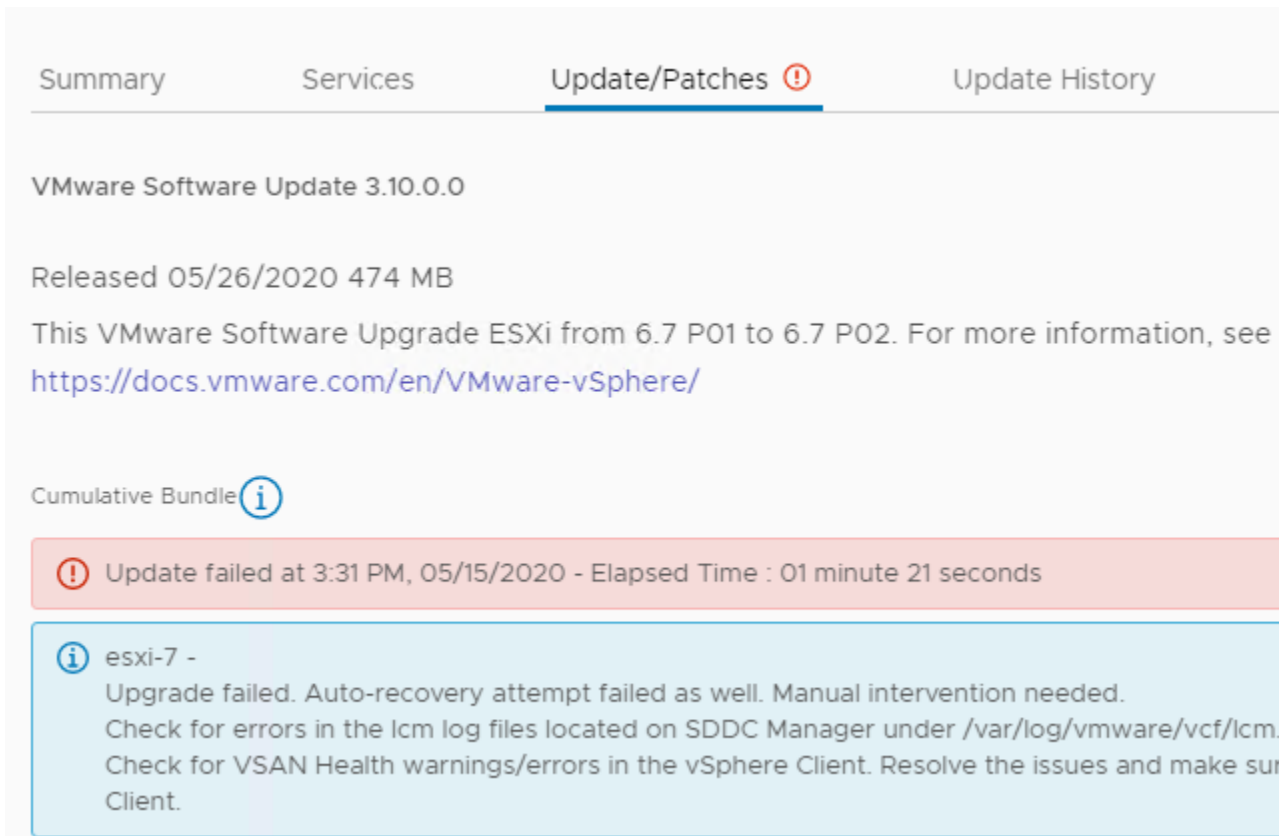
Released 06/23/2020 11 GB

This VMware Cloud Foundation Upgrade 4.0.0.1 to 4.0.1.0 contains features, critical bugs and security f
<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.01/rn/VMware-Cloud-Foundation-401-Release>

Updating OPERATIONS MANAGER

<div style="display: flex; align-items: center;"> } SDDC MANAGER </div>	○ In Progress
<div style="display: flex; align-items: center;"> > COMMON SERVICES </div>	✓ Updated 🔄
<div style="display: flex; align-items: center;"> > OPERATIONS MANAGER </div>	○ In Progress
<div style="display: flex; align-items: center;"> ⌚ DOMAIN MANAGER </div>	⌚ Queued
<div style="display: flex; align-items: center;"> ⌚ SOLUTIONS MANAGER </div>	⌚ Queued
<div style="display: flex; align-items: center;"> ⌚ SDDC MANAGER UI </div>	⌚ Queued
<div style="display: flex; align-items: center;"> ⌚ LCM </div>	⌚ Queued
<div style="display: flex; align-items: center;"> ⌚ MULTI SITE SERVICE </div>	⌚ Queued

- When all tasks to update a component have been completed, the update status for the component is displayed as Updated.
- If a component fails to be updated, the status is displayed as Failed. The reason for the failure as well as remediation steps are displayed. The image below is an example and may not reflect the actual versions in your environment.



Summary Services **Update/Patches** Update History

VMware Software Update 3.10.0.0

Released 05/26/2020 474 MB

This VMware Software Upgrade ESXi from 6.7 P01 to 6.7 P02. For more information, see <https://docs.vmware.com/en/VMware-vSphere/>

Cumulative Bundle

! Update failed at 3:31 PM, 05/15/2020 - Elapsed Time : 01 minute 21 seconds

i esxi-7 -
Upgrade failed. Auto-recovery attempt failed as well. Manual intervention needed.
Check for errors in the lcm log files located on SDDC Manager under /var/log/vmware/vcf/lcm.
Check for VSAN Health warnings/errors in the vSphere Client. Resolve the issues and make sur Client.

6. After you resolve the issues, you can retry the update.

View VMware Cloud Foundation Update History

The Update History page displays all updates applied to a workload domain.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. Click the name of a workload domain and then click the **Update History** tab.
All updates applied to this workload domain are displayed. If an update bundle was applied more than once, click **View Past Attempts** to see more information.

Access VMware Cloud Foundation Upgrade Log Files

You can check the log files for failed upgrades to help troubleshoot and resolve issues.

1. SSH in to the SDDC Manager appliance with the `vcf` user name and enter the password.
2. To access upgrade logs, navigate to the `/var/log/vmware/vcf/lcm` directory.
 - `lcm-debug` log file contains debug level logging information.
 - `lcm.log` contains information level logging.
3. To create an sos bundle for support, see Supportability and Serviceability (SoS) Utility in the *VMware Cloud Foundation Administration Guide*.

Downloading VMware Cloud Foundation Upgrade Bundles

Before you can upgrade VMware Cloud Foundation, you must download the upgrade bundles for each VMware Cloud Foundation component that requires an upgrade.

Online and Offline Downloads

If the SDDC Manager appliance can connect to the internet (either directly or through a proxy server), you can download upgrade bundles from the VMware Depot using your Broadcom Support Portal account.

If the SDDC Manager appliance cannot connect to the internet, you can use the Bundle Transfer Utility or connect to an offline depot.

See [Public URL list for SDDC Manager](#) for information about the URLs that must be accessible to download bundles.

Other Bundle Types

In addition to upgrade bundles, VMware Cloud Foundation includes the following bundle types:

- **Install Bundles**
An install bundle includes software binaries to install VI workload domains (vCenter Server and NSX) and VMware Aria Suite Lifecycle. You download install bundles using the same process that you use for upgrade bundles.
- **Async Patch Bundles**
An async patch bundle allows you to apply critical patches to certain VMware Cloud Foundation components (NSX Manager, vCenter Server, and ESXi) when an update or upgrade bundle is not available. If you are running VMware Cloud Foundation 5.1 or earlier, you must use the Async Patch Tool to download an async patch bundle. See [Async Patch Tool](#). Starting with VMware Cloud Foundation 5.2, you can download async patches using the SDDC Manager UI or Bundle Transfer Utility.

Connect SDDC Manager to a Software Depot for Downloading Bundles

SDDC Manager can connect to a software depot to download software bundles, compatibility data, and more.

To connect to the online depot, SDDC Manager must be able to connect to the internet, either directly or through a proxy server.

To connect to an offline depot, you must first configure it. See [KB 312168](#) for information about the requirements and process for creating an offline depot. To download bundles to an offline depot, see "Download Bundles to an Offline Depot" in the *VMware Cloud Foundation Lifecycle Management Guide*.

SDDC Manager supports two types of software depots:


- Online depot
- Offline depot

You can only connect SDDC Manager to one type of depot. If SDDC Manager is connected to an online depot and you configure a connection to an offline depot, the online depot connection is disabled and deleted.

1. In the navigation pane, click **Administration** > **Depot Settings**.

Online Depots

VMware Depot

 Depot connection not set up. Authenticate your Customer Connect Account to set it up.

AUTHENTICATE

Offline Depot

Offline Depot

 Depot connection not set up.

SET UP

2. Connect SDDC Manager to an online depot or an offline depot.

Depot Type	Configuration Steps
Online	<ol style="list-style-type: none"> 1. Click Authenticate for the VMware Depot. 2. Type your Broadcom Support Portal user name and password. 3. Click Authenticate
Offline	<ol style="list-style-type: none"> 1. Click Set Up for the Offline Depot. 2. Enter the following information for the offline depot: <ul style="list-style-type: none"> – FQDN or IP address – Port number – User name – Password 3. Click Set Up.

SDDC Manager attempts to connect to the depot. If the connection is successful, SDDC Manager starts looking for available bundles. To view available bundles, click **Lifecycle Management > Bundle Management** and then click the **Bundles** tab. It may take some time for all available bundles to appear.

Download Bundles Using SDDC Manager

After you connect SDDC Manager to an online or offline depot, you can view and download available upgrade bundles.

Connect SDDC Manager to an online or offline depot. See [Connect SDDC Manager to a Software Depot for Downloading Bundles](#).

If SDDC Manager does not have direct internet access, configure a proxy server or use the Bundle Transfer Utility for offline bundle downloads.

- [Configure a Proxy Server for Downloading Bundles](#)
- [Offline Bundle Download for VMware Cloud Foundation](#)

When you download bundles, SDDC Manager verifies that the file size and checksum of the downloaded bundles match the expected values.

1. In the navigation pane, click **Lifecycle Management** > **Bundle Management**.
2. Click the **Bundles** tab to view available bundles.

NOTE

If you just connected SDDC Manager to a depot, it can take some time for bundles to appear.

All available bundles are displayed. Install bundles display an Install Only Bundle label. If the bundle can be applied right away, the Bundle Details column displays the workload domains to which the bundle needs to be applied to, and the Availability column says Available. If another bundle must be applied before a particular bundle, the Availability field displays Future.

To view more information about the bundle, click **View Details**. The Bundle Details section displays the bundle version, release date, and additional details about the bundle.

3. For the bundle you want to download, do one of the following:

- Click **Download Now** for an immediate download.

The bundle download begins right away.

- Click **Schedule Download** to schedule a download.

Select the date and time for the bundle download and click **Schedule**.

4. Click the **Download History** tab to see the downloaded bundles.

Configure a Proxy Server for Downloading VMware Cloud Foundation Bundles

If SDDC Manager does not have direct internet access, you can configure a proxy server to download bundles. VMware Cloud Foundation 5.2 and later support proxy servers with authentication.

1. In the navigation pane, click **Administration** > **Proxy Settings**.
2. Click **Set Up Proxy**.
3. Toggle the **Enable Proxy** setting to the on position.
4. Select **HTTP** or **HTTPS**.
5. Enter the proxy server IP address and port number.
6. If your proxy server requires authentication, toggle the **Authentication** setting to the on position and enter the user name and password.
7. Click **Save**.

You can now download bundles as described in [Download Bundles Using SDDC Manager](#).

Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles

If the SDDC Manager appliance does not have access to the VMware Depot, you can use the Bundle Transfer Utility to download the bundles to a different computer and then upload them to the SDDC Manager appliance.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.

NOTE

The Bundle Transfer Utility is the only supported method for downloading bundles. Do not use third-party tools or other methods to download bundles.

Using the Bundle Transfer Utility to upgrade to VMware Cloud Foundation 5.2.x involves the following steps:

- Download the latest version of the Bundle Transfer Utility.
- On a computer with access to the internet, use the Bundle Transfer Utility to download the bundles and other required files.
- Copy the bundles and other required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and other required files to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS. To use this option, the proxy certificate must be added to Bundle Transfer Utility JRE default trust store. For example: <code>/opt/obtu/jre/lin64/bin/keytool -importcert -file ca-bundle.crt -keystore /opt/obtu/jre/lin64/lib/security/cacerts</code>
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUserUsername--proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../password.txt --proxyHttps
```

1. Download the most recent version of the Bundle Transfer Utility on a computer with internet access.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
 - b) Click the version of VMware Cloud Foundation to which you are upgrading.
 - c) Click **Drivers & Tools**.
 - d) Click the download icon for the Bundle Transfer Utility.
 - e) Extract `lcm-tools-prod.tar.gz`.
 - f) Navigate to the `lcm-tools-prod/bin/` and confirm that you have execute permission on all folders.

2. Download bundles and other artifacts to the computer with internet access.

- a) Download the manifest file.

This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUserUsername
```

For `--depotUser`, enter your Broadcom Support Portal user name.

Note the location to which the Bundle Transfer Utility downloads the manifest. You will use this as the `--sourceManifestDirectory` when you upload the manifest. For example:

```
Validating the depot user credentials...
Downloading LCM Manifest to: /root/PROD2/evo/vmw
Successfully completed downloading file
```

- b) Download the compatibility data.

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUserUsername
```

To specify a download location, use `--outputDirectory` followed by the path to the directory.

- c) Download the vSAN HCL file.

```
./lcm-bundle-transfer-util --vsanHclDownload
```

- d) Download the upgrade bundles.

```
./lcm-bundle-transfer-util --download --outputDirectoryabsolute-path-output-dir-
--depotUserUsername--svcurrent-vcf-version--ptarget-vcf-version
```

where

<code>absolute-path-output-dir</code>	Path to the directory where the bundle files should be downloaded. This directory folder must have <code>777</code> permissions. If you do not specify the download directory, bundles are downloaded to the default directory with <code>777</code> permissions.
<code>depotUser</code>	User name for the Broadcom Support Portal. You will be prompted to enter the user password. If there are any special characters in the password, specify the password within single quotes.
<code>current-vcf-version</code>	Current version of VMware Cloud Foundation. For example, <code>4.5.2.0</code> .
<code>target-vcf-version</code>	Target version of VMware Cloud Foundation. For example, <code>5.2.1.0</code> .

Follow the prompts in the Bundle Transfer Utility.

- e) Specify the bundles to download.

Enter one of the following options:

- all
- install
- patch

You can also enter a comma-separated list of bundle names to download specific bundles. For example: `bundle-38371, bundle-38378`.

Download progress for each bundle is displayed. Wait until all bundles are downloaded successfully.

3. Copy the following files/directories to the SDDC Manager appliance.

- Bundle Transfer Utility
- Manifest file
- Compatibility data file (`VmwareCompatibilityData.json`)
- vSAN HCL
- Entire bundle output directory

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

4. Copy the bundle transfer utility to the SDDC Manager appliance.

- a) SSH in to the SDDC Manager appliance using the `vcf` user account.
- b) Enter `su` to switch to the root user.
- c) Create the `lcm-tools` directory.

```
mkdir /opt/vmware/vcf/lcm/lcm-tools
```

NOTE

If the `/opt/vmware/vcf/lcm/lcm-tools` directory already exists with an older version of the Bundle Transfer Utility, you need to delete contents of the existing directory before proceeding.

- d) Copy the Bundle Transfer Utility file (`lcm-tools-prod.tar.gz`) that you downloaded in step 1 to the `/opt/vmware/vcf/lcm/lcm-tools` directory.
- e) Extract the contents of `lcm-tools-prod.tar.gz`.
- f) Set the permissions for the `lcm-tools` directory.

```
cd /opt/vmware/vcf/lcm/
chown vcf_lcm:vcf -R lcm-tools
chmod 750 -R lcm-tools
```

5. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and artifacts.

- a) Upload the manifest file.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectoryManifest-Directory-  
-sddcMgrFqdn FQDN--sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

- b) Upload the compatibility file.

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --
inputDirectorycompatibility-file-directory--sddcMgrFqdnFQDN--sddcMgrUserUsernam
e
```

c) Upload the HCL file.

```
./lcm-bundle-transfer-util --vsanHclUpload --inputDirectoryhcl-file-path--
sddcMgrFqdn.sddc-manager-fqdn--sddcMgrUseruser
```

d) Upload the bundle directory.

```
./lcm-bundle-transfer-util --upload --bundleDirectoryabsolute-path-bundle-dir
```

Offline Download of Independent SDDC Manager Bundles

Once SDDC Manager is upgraded to 5.2 or later, new functionality is introduced that allows you to get the latest SDDC Manager features and security fixes without having to upgrade the entire VMware Cloud Foundation BOM. This procedure describes using the Bundle Transfer Utility to download an SDDC Manager bundle released independently of the VMware Cloud Foundation BOM when SDDC Manager is not connected to an online depot..

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- The computer with internet connectivity and the SDDC Manager appliance must have the latest version of the Bundle Transfer Utility installed and configured. See [Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles](#) for more information.

An independent SDDC Manager release includes a fourth digit in its version number, for example SDDC Manager 5.2.0.1.

- On a computer with access to the internet, use the Bundle Transfer Utility to download the independent SDDC Manager bundle and other required files.
- Copy the bundle and other required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundle and other required files to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
--proxyServer, --ps	Provide the proxy server FQDN and port. For example: --proxyServer proxy.example.com:3128.
--proxyHttps	Add this option if the proxy server uses HTTPS.
--proxyUser	For a proxy server that requires authentication, enter the user name.
--proxyPasswordFile	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, --proxyPasswordFile ../../password.txt.

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download bundles and other artifacts to the computer with internet access.

a) Download the manifest file.

This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username
```

For `--depotUser`, enter your Broadcom Support Portal user name.

b) Download the compatibility data.

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUser
Username
```

To specify a download location, use `--outputDirectory` followed by the path to the directory.

c) Download the independent SDDC Manager upgrade bundle.

```
./lcm-bundle-transfer-util --download --sddcMgrVersion four-digit-sddc-version
--depotUser Username --outputDirectory absolute-path-output-dir
```

where

<code>depotUser</code>	User name for the Broadcom Support Portal. You will be prompted to enter the user password. If there are any special characters in the password, specify the password within single quotes.
<code>four-digit-sddc-version</code>	Target version of SDDC Manager. For example, 5.2.0.1.
<code>absolute-path-output-dir</code>	Path to the directory where the bundle files should be downloaded. This directory folder must have 777 permissions. If you do not specify the download directory, bundles are downloaded to the default directory with 777 permissions.

Follow the prompts in the Bundle Transfer Utility.

2. Copy the following files/directories to the SDDC Manager appliance.

- Manifest file
- Compatibility data file (`VmwareCompatibilityData.json`)
- Entire bundle output directory

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

3. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and artifacts.

a) Upload the manifest file.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory Manifest-
Directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

b) Upload the compatibility file.


```
./lcm-bundle-transfer-util --update --compatibilityMatrix --inputDirectory
compatibility-file-directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

c) Upload the bundle directory.

```
./lcm-bundle-transfer-util --upload --bundleDirectory absolute-path-bundle-dir
```

After the upload completes successfully, you can use the SDDC Manager UI to upgrade SDDC Manager. See [Independent SDDC Manager Upgrade using the SDDC Manager UI](#).

Offline Download of Async Patch Bundles

Once SDDC Manager is upgraded to 5.2 or later, a new option for patching VMware Cloud Foundation components is available in the SDDC Manager UI. This procedure describes using the Bundle Transfer Utility to download async patches when SDDC Manager is not connected to an online depot.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- The computer with internet connectivity and the SDDC Manager appliance must have the latest version of the Bundle Transfer Utility installed and configured. See [Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles](#) for more information.

Offline download of async patches involves the following steps:

- On a computer with access to the internet, use the Bundle Transfer Utility to download the async patch bundle and other required files.
- Copy the bundle and other required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundle and other required files to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download bundles and other artifacts to the computer with internet access.

a) Download the manifest file.

This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username
```

For `--depotUser`, enter your Broadcom Support Portal user name.

b) Download the compatibility data.

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUser Username
```

To specify a download location, use `--outputDirectory` followed by the path to the directory.

c) Download the product version catalog.

```
./lcm-bundle-transfer-util --depotUser Username --download productVersionCatalog --outputDirectory directory-path
```

d) List the available async patches.

```
./lcm-bundle-transfer-util --listAsyncPatchBundles --depotUser Username
```

e) Download an async patch.

```
./lcm-bundle-transfer-util --download --bundle bundle-number --depotUser Username
```

For example:

```
./lcm-bundle-transfer-util --download --bundle bundle-12345 --depotUser user@example.com
```

2. Copy the following files/directories to the SDDC Manager appliance.

- Manifest file
- Compatibility data file (`VmwareCompatibilityData.json`)
- Entire bundle output directory

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

3. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and artifacts.

a) Upload the manifest file.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory Manifest-Directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

b) Upload the compatibility file.

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --inputDirectory compatibility-file-directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

c) Upload the product version catalog.

```
./lcm-bundle-transfer-util --upload productVersionCatalog --inputDirectory directory-path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

d) Upload the bundle directory.

```
./lcm-bundle-transfer-util --upload --bundle bundle-number --bundleDirectory
absolute-path-bundle-dir
```

- Replace *number* with the bundle number you are uploading. For example: 12345 for bundle-12345.
- Replace *absolute-path-bundle-dir* with the path to the location where you copied the output directory. For example: /nfs/vmware/vcf/nfs-mount/upgrade-bundles.

After the upload completes successfully, you can use the SDDC Manager UI to apply the async patch. See [Patching the Management and Workload Domains](#).

Offline Download of Flexible BOM Upgrade Bundles

Once SDDC Manager is upgraded to version 5.2 or later, new functionality is introduced to the upgrade planner that allows you to select specific target versions for each VMware Cloud Foundation component you want to upgrade. This procedure describes using the Bundle Transfer Utility to download the bundles for a flexible BOM upgrade when SDDC Manager is not connected to an online depot.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- A Windows or Linux computer with access to the SDDC Manager appliance for uploading the bundles.
- To upload the manifest file from a Windows computer, you must have OpenSSL installed and configured.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- The computer with internet connectivity and the SDDC Manager appliance must all have the latest version of the Bundle Transfer Utility installed and configured. See [Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles](#) for more information.

After you download the bundles, you can use the upgrade planner in the SDDC Manager UI to select any supported version for each of the VMware Cloud Foundation BOM components. This includes async patch versions as well as VCF BOM versions.

Offline download of flexible BOM upgrade bundles involves the following steps:

- On a computer with access to the internet, use the Bundle Transfer Utility to download the required files.
- Copy the required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the required files to the internal LCM repository.
- Plan the upgrade using the SDDC Manager UI.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to generate the `plannerFile.json`.
- Copy `plannerFile.json` to the computer with internet access.
- On the computer with access to the internet, download bundles using `plannerFile.json`.
- Copy the bundle directory to the SDDC Manager appliance and use the Bundle Transfer Utility to upload the bundles to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.

Table continued on next page

Continued from previous page

Option	Description
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download the required files to the computer with internet access.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
outputDirectory directory-path
```

The manifest is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

For `--depotUser`, enter your Broadcom Support Portal user name.

```
./lcm-bundle-transfer-util --download --bundleManifests --depotUser Username --
bundleManifestsDir directory-path
```

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUser Username --
pdu dell_depot_email --outputDirectory directory-path
```

```
./lcm-bundle-transfer-util --depotUser Username --download productVersionCatalog --
outputDirectory directory-path
```

2. Copy the entire output directory to the SDDC Manager appliance.

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

3. On the SDDC Manager appliance, upload/update the files.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory directory-path --
sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

```
./lcm-bundle-transfer-util --upload --bundleManifests --bundleManifestsDir directory-
path
```

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --inputDirectory directory-
path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

```
./lcm-bundle-transfer-util --upload productVersionCatalog --inputDirectory directory-
path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

- In the SDDC Manager UI, plan the upgrade.

See [Flexible BOM Upgrade in VMware Cloud Foundation](#).

- On the SDDC Manager appliance, use the Bundle Transfer Utility to generate a planner file.

```
./lcm-bundle-transfer-util --generatePlannerFile --sddcMgrUser Username --sddcMgrFqdn FQDN --outputDirectory directory-path --domainNames domain-name -p target-vcf-version
```

For example:

```
./lcm-bundle-transfer-util --generatePlannerFile --sddcMgrUser administrator@vsphere.local --sddcMgrFqdn sddc-manager.example.com --outputDirectory /home/vcd --domainNames mgmt-domain -p 5.2.0.0
```

- Copy `plannerFile.json` file to the computer with access to the internet.
- On the computer with access to the internet, download the bundles using the `plannerFile.json`.

```
./lcm-bundle-transfer-util --download --plannerFile directory-path --depotUser Username
```

- Copy the entire output directory to the SDDC Manager appliance.
- Upload the bundle directory to the SDDC Manager appliance internal LCM repository.

```
./lcm-bundle-transfer-util --upload --bundleDirectory directory-path
```

In the SDDC Manager UI browse to the Available Updates screen for the workload domain you are upgrading and click **Schedule Update** or **Update Now** to update the first component. Continue to update the VCF BOM components until they are all updated.

HCL Offline Download for VMware Cloud Foundation

If the SDDC Manager appliance does not have access to the VMware Depot, you can use the Bundle Transfer Utility to manually download the HCL file from the depot on your local computer and then upload it to the SDDC Manager appliance.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the HCL. To upload the HCL file from a Windows computer, you must have OpenSSL installed and configured.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.

NOTE

The Bundle Transfer Utility is the only supported method for downloading HCL. Do not use third-party tools or other methods to download HCL.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading the HCL:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.

Table continued on next page

Continued from previous page

Option	Description
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --vsanHclDownload --outputDirectory output-directory --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download the most recent version of the Bundle Transfer Utility on a computer with internet access.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
 - b) Click the version of VMware Cloud Foundation to which you are upgrading.
 - c) Click **Drivers & Tools**.
 - d) Click the download icon for the Bundle Transfer Utility.
2. Extract `lcm-tools-prod.tar.gz`.
3. Navigate to the `lcm-tools-prod/bin/` and confirm that you have execute permission on all folders.
4. Copy the bundle transfer utility to a computer with access to the SDDC Manager appliance and then copy the bundle transfer utility to the SDDC Manager appliance.
 - a) SSH in to the SDDC Manager appliance using the `vcf` user account.
 - b) Enter `su` to switch to the root user.
 - c) Create the `lcm-tools` directory.

```
mkdir /opt/vmware/vcf/lcm/lcm-tools
```

NOTE

If the `/opt/vmware/vcf/lcm/lcm-tools` directory already exists with an older version of the Bundle Transfer Utility, you need to delete contents of the existing directory before proceeding.

- d) Copy the Bundle Transfer Utility file (`lcm-tools-prod.tar.gz`) that you downloaded in step 1 to the `/opt/vmware/vcf/lcm/lcm-tools` directory.
- e) Extract the contents of `lcm-tools-prod.tar.gz`.

```
tar -xvf lcm-tools-prod.tar.gz
```

- f) Set the permissions for the `lcm-tools` directory.

```
cd /opt/vmware/vcf/lcm/
chown vcf_lcm:vcf -R lcm-tools
chmod 750 -R lcm-tools
```

5. On the computer with internet access, download the HCL file.

```
./lcm-bundle-transfer-util --vsanHclDownload --outputDirectory output-directory
```

It can also be downloaded to the default path:

```
./lcm-bundle-transfer-util --vsanHclDownload
```

6. Copy the HCL file to the SDDC Manager appliance.
7. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the HCL file.

```
./lcm-bundle-transfer-util --vsanHclUpload --inputDirectory hcl-file-path --
sddcMgrFqdn sddc-manager-fqdn --sddcMgrUser user
```

hcl-file-path	Path from where HCL file should be picked up to upload. e.g /root/testdownload/vsan/hcl/all.json. If not given default will be taken. (/root/PROD2/vsan/hcl/all.json)
sddc-manager-fqdn	SDDC Manager FQDN. If not given default will be taken.
user	SDDC Manager user. After this, the tool will prompt for the user password.

Download Bundles to an Offline Depot

VMware Cloud Foundation 5.2 and later support an offline depot that you can connect to from multiple instances of SDDC Manager. Use the Bundle Transfer Utility to download and transfer bundles to the offline depot and then any SDDC Manager connected to the offline depot can access the bundles.

- [Set up an offline depot.](#)
- The offline depot must have:
 - The latest version of the Bundle Transfer Utility. You can download it from the Broadcom Support portal.
 - Internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
 - Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- Connect SDDC Manager to the offline depot. See [Connect SDDC Manager to a Software Depot for Downloading Bundles](#).

NOTE

You can also connect SDDC Manager to the offline depot after you download bundles to the offline depot.

You can use the Bundle Transfer Utility to download upgrade bundles and async patch bundles to the offline depot.

1. On the computer hosting offline depot, run the following command to download the bundles required to upgrade VMware Cloud Foundation.

```
./lcm-bundle-transfer-util --setUpOfflineDepot -sv vcf-source-version --
offlineDepotRootDir offline-depot-root-dir --offlineDepotUrl url:port --depotUser
user-name --depotUserPasswordFile path-to-password-file
```

For example:

```
./lcm-bundle-transfer-util --setUpOfflineDepot -sv 5.0.0.0 --offlineDepotRootDir /
var/www --offlineDepotUrl https://10.123.456.78:8282 --depotUser user@example.com --
depotUserPasswordFile ../vmw-depot
```

2. Run the following command to download async patch bundles to the offline depot:

```
./lcm-bundle-transfer-util --setUpOfflineDepot --asyncPatches -offlineDepotRootDir
offline-depot-root-dir --offlineDepotUrl url:port --depotUser user-name --
depotUserPasswordFile path-to-password-file
```

For example:

```
./lcm-bundle-transfer-util --setUpOfflineDepot --asyncPatches -offlineDepotRootDir /
var/www --offlineDepotUrl https://10.123.456.78:8282 --depotUser user@example.com --
depotUserPasswordFile ../vmw-depot
```

After the bundles are available in the offline depot, you can use the SDDC Manager UI to apply the bundles to workload domains. Multiple instances of SDDC Manager UI can connect to the same offline depot.

VMware Cloud Foundation Upgrade Prerequisites

Before you upgrade VMware Cloud Foundation, make sure that the following prerequisites are met.

Table 202: Upgrade Prerequisites

Prerequisite	Additional Information
Allocate a temporary IP address for each vCenter Server upgrade	[Conditional] When upgrading from VMware Cloud Foundation 4.5.x. Required for each vCenter Server upgrade. Must be allocated from the management subnet. The IP address can be reused.
Obtain updated licenses	New licenses required for: <ul style="list-style-type: none"> • vSAN 8.x • vSphere 8.x
Verify there are no expired or expiring passwords	Review the password management dashboard in SDDC Manager.
Verify there are no expired or expiring certificates	Review the Certificates tab in SDDC Manager for each workload domain.
Verify ESXi host TPM module status	[Conditional] If ESXi hosts have TPM modules in use, verify they are running the latest 2.0 firmware. If not in use they must be disabled in the BIOS. See KB 312159
Verify ESXi hardware is compatible with target version	See ESXi Requirements and VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php .
Manually update the vSAN HCL database to ensure that it is up-to-date.	See KB 2145116
Back up SDDC Manager, all vCenter Server instances, and NSX Manager instances.	Take file-based backups or image-level backups of SDDC Manager, all vCenter Servers, and NSX Managers. Take a cold snapshot of SDDC Manager.
Make sure that there are no failed workflows in your system and none of the VMware Cloud Foundation resources are in activating or error state.	CAUTION If any of these conditions are true, contact VMware Technical Support before starting the upgrade.
Review the <i>Release Notes</i> for known issues related to upgrades.	
Deactivate all VMware Cloud Foundation 4.x async patches and run an inventory sync before upgrading.	VMware Cloud Foundation 5.0 and later no longer require using the Async Patch Tool to enable upgrades from an

Table continued on next page

Continued from previous page

Prerequisite	Additional Information
	async-patched VMware Cloud Foundation instance. See VCF Async Patch Tool Options for more information
Review Operational Impacts of NSX Upgrade in <i>NSX Upgrade Guide</i> to understand the impact that each component upgrade might have on your environment.	
In the vSphere Client, ensure there are no active alarms on hosts or vSphere clusters.	
Download the upgrade bundles.	See Downloading VMware Cloud Foundation Upgrade Bundles .

VMware Cloud Foundation 5.2.x Upgrade Overview

This section describes the tasks required to perform an upgrade to VMware Cloud Foundation 5.2.x.

VMware Cloud Foundation Upgrade Preparation

Review the [VMware Cloud Foundation Upgrade Prerequisites](#) before starting an upgrade.

Management Domain Upgrade

Table 203: SDDC Manager Upgrade

Task	Applies When	Additional Information
<ul style="list-style-type: none"> Precheck Update - Versions Prior to SDDC Manager 5.0 Perform Update Precheck in SDDC Manager 		
Apply the VMware Cloud Foundation Upgrade Bundle	<ul style="list-style-type: none"> The initial VMware Cloud Foundation version is <ul style="list-style-type: none"> – 4.5.x or 5.x 	If the current version of VMware Cloud Foundation is 4.5.x or 5.x Upgrade SDDC Manager to 5.2.x.
Apply the VMware Cloud Foundation Configuration Updates	<ul style="list-style-type: none"> Once the SDDC Manager has been upgraded to 5.2.x the Configuration updates can be applied collectively. 	
Update Compatibility Data with the Bundle Transfer Utility		[Conditional] Required when using offline bundle download

Table 204: Upgrade VMware Aria Suite

Task	Additional Information
Upgrade VMware Aria Suite Lifecycle for VMware Cloud Foundation	[Conditional] If VMware Aria Suite Lifecycle is present
Upgrade VMware Aria Suite products for VMware Cloud Foundation	[Conditional] If VMware Aria Suite products are present

Table 205: Upgrade NSX With Federation

Task	Applies When	Additional Information
Upgrade NSX Global Managers to 4.2	When NSX is deployed in the workload domain with NSX Federation configured.	<ul style="list-style-type: none"> • [Conditional] If NSX Federation is present • Upgrade NSX Global Managers to 4.2 using the Global Manager UI • Upgrade standby global manager, followed by active global manager • [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade all NSX global managers in your estate now.
Upgrade to NSX 4.2		<ul style="list-style-type: none"> • Upgrade NSX to 4.2 using SDDC Manager • [Optional] If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now. • NSX upgrades across VI workload domains can be completed in sequence or up to five in parallel.

Table 206: Upgrade NSX Without Federation

Task	Applies When	Additional Information
Upgrade to NSX 4.2	When NSX is deployed in the workload domain and is not using NSX Federation.	<ul style="list-style-type: none"> • Upgrade NSX to 4.2 using SDDC Manager. • [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now.

Table 207: Upgrade vCenter Server

Task	Additional Information
Upgrade vCenter Server for VMware Cloud Foundation	<ul style="list-style-type: none"> • [Conditional] When upgrading from VMware Cloud Foundation 4.5.x. Requires a temporary IP address in the management subnet • [Conditional] When upgrading to VMware Cloud Foundation 5.2.1 using vCenter Reduced Downtime Upgrade (RDU).

Table continued on next page

Continued from previous page

Task	Additional Information
	<p>Requires a temporary IP address in the management subnet</p> <ul style="list-style-type: none"> • [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade vCenter Servers that share a SSO Domain across all VI workload domains now in a serial order. Isolated Workload Domains can be upgraded in parallel

Table 208: Upgrade Management Domain vSphere clusters

Task	Additional Information
Upgrade vSAN Witness Host for VMware Cloud Foundation	[Conditional] If the vSphere cluster is a stretched vSAN cluster
Skip Hosts During vSphere clusters Update	[Conditional] If you need to skip hosts
Upgrade vSphere clusters with Custom ISOs or Upgrade vSphere clusters with VMware Cloud Foundation Stock ISO and Async Drivers or Upgrade vSphere clusters with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation or Upgrade vSphere clusters with vSphere Lifecycle Manager Images for VMware Cloud Foundation	<ul style="list-style-type: none"> • Choose an approach based on your requirements. • [Optional] If you are upgrading by component rather than by workload domain, upgrade vSphere clusters across all VI workload domains now.

Table 209: Post Upgrade Tasks

Task	Additional Information
Update Licenses for a Workload Domain	<p>[Conditional] If upgrading from a VMware Cloud Foundation version prior to 5.0</p> <p>Update licenses for:</p> <ul style="list-style-type: none"> • vSAN 8.x • vSphere 8.x
Apply Configuration Updates	[Conditional] If there are configuration updates required
Upgrade vSphere Distributed Switch versions	<ul style="list-style-type: none"> • [Optional] The upgrade lets the distributed switch take advantage of features that are available only in the later versions.

Table continued on next page

Continued from previous page

Task	Additional Information
Upgrade vSAN on-disk format versions	<ul style="list-style-type: none"> • The upgrade lets the vSAN Cluster take advantage of features that are available only in the later versions. • The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure. • These updates can be performed at a time that is most convenient for your organization..

VI Workload Domain Upgrade

Table 210: Upgrade Precheck

Task	Additional Information
Perform an upgrade precheck	

Table 211: Upgrade NSX Without Federation

Task	Applies When	Additional Information
Upgrade to NSX 4.2	When NSX is deployed in the workload domain and is not using NSX Federation.	<ul style="list-style-type: none"> • Upgrade NSX to 4.2 using SDDC Manager. • [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now.

Table 212: Upgrade NSX With Federation

Task	Applies When	Additional Information
Upgrade NSX Global Managers to 4.2	When NSX is deployed in the workload domain with NSX Federation configured.	<ul style="list-style-type: none"> • [Conditional] If NSX Federation is present • Upgrade NSX Global Managers to 4.2 using the Global Manager UI • Upgrade standby global manager, followed by active global manager • [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade all NSX global managers in your estate now.
Upgrade to NSX 4.2		<ul style="list-style-type: none"> • Upgrade NSX to 4.2 using SDDC Manager

Table continued on next page

Continued from previous page

Task	Applies When	Additional Information
		<ul style="list-style-type: none"> • [Optional] If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now. • NSX upgrades across VI workload domains can be completed in sequence or up to five in parallel.

Table 213: Upgrade vCenter Server

Task	Additional Information
Upgrade vCenter Server for VMware Cloud Foundation	<ul style="list-style-type: none"> • [Conditional] When upgrading from VMware Cloud Foundation 4.5.x. Requires a temporary IP address in the management subnet • [Conditional] When upgrading to VMware Cloud Foundation 5.2.1 using vCenter Reduced Downtime Upgrade (RDU). Requires a temporary IP address in the management subnet • [Conditional] for VI Workload Domain upgrades. If you are upgrading by component rather than by workload domain, upgrade vCenter Servers that share a SSO Domain across all VI workload domains now in a serial order. Isolated Workload Domains can be upgraded in parallel

Table 214: Upgrade VI Workload Domain vSphere clusters

Task	Additional Information
Upgrade vSAN Witness Host for VMware Cloud Foundation	[Conditional] If the vSphere cluster is a stretched vSAN cluster
Skip Hosts During vSphere clusters Update	[Conditional] If you need to skip hosts
Upgrade vSphere clusters with Custom ISOs or Upgrade vSphere clusters with VMware Cloud Foundation Stock ISO and Async Drivers or Upgrade vSphere clusters with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation or	<ul style="list-style-type: none"> • Choose an approach based on your requirements • [Optional] If you are upgrading by component rather than by workload domain, upgrade vSphere clusters across all VI workload domains now.

Table continued on next page

Continued from previous page

Task	Additional Information
Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation	
Post Upgrade Steps for NFS-Based VI Workload Domains	

Table 215: Post Upgrade Tasks

Task	Additional Information
Update Licenses for a Workload Domain	<p>[Conditional] If upgrading from a VMware Cloud Foundation version prior to 5.0</p> <p>Update licenses for:</p> <ul style="list-style-type: none"> • vSAN 8.x • vSphere 8.x
Apply Configuration Updates	[Conditional] If there are configuration updates required
Upgrade vSphere Distributed Switch versions	<ul style="list-style-type: none"> • [Optional] The upgrade lets the distributed switch take advantage of features that are available only in the later versions.
Upgrade vSAN on-disk format versions	<ul style="list-style-type: none"> • The upgrade lets the vSAN Cluster take advantage of features that are available only in the later versions. • The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure. • These updates can be performed at a time that is most convenient for your organization..

Upgrade the Management Domain to VMware Cloud Foundation 5.2.x

To upgrade to VMware Cloud Foundation 5.2.x, the management domain must be at VMware Cloud Foundation 4.5 or higher. If your environment is at a version lower than 4.5, you must upgrade the management domain to 4.5 or later and then upgrade to 5.2.x.

Until SDDC Manager is upgraded to version 5.2.x, you must upgrade the management domain before you upgrade VI workload domains. Once SDDC Manager is at version 5.2 or later, you can upgrade VI workload domains before or after upgrading the management domain, as long as all components in the workload domain are compatible.

Upgrade the components in the management domain in the following order:

1. SDDC Manager and VMware Cloud Foundation services.
2. VMware Aria Suite Lifecycle
3. NSX Manager and NSX Global Managers (if applicable).
4. vCenter Server.
5. ESXi

After all upgrades have completed successfully:

1. Remove the VM snapshots you took before starting the update.
2. Take a backup of the newly installed components.


Perform Update Precheck - Versions Prior to SDDC Manager 5.0

If you have not yet upgraded to SDDC Manager 5.0, these are the steps to run a Precheck. You must perform a precheck before applying an update or upgrade bundle to ensure that your environment is ready for the update.

For an ESXi bundle, the system performs a bundle level precheck in addition to the environment precheck. For VI workload domains using vSphere Lifecycle Manager baselines, the ESXi bundle precheck validates the following.

- Custom ISO is compatible with your environment.
- Custom ISO size is smaller than the boot partition size.
- Third party VIBs are compatible with the environment.

If you silence a vSAN Skyline Health alert in the vSphere Client, SDDC Manager skips the related precheck and indicates which precheck it skipped. Click **Restore Precheck** to include the silenced precheck. For example:

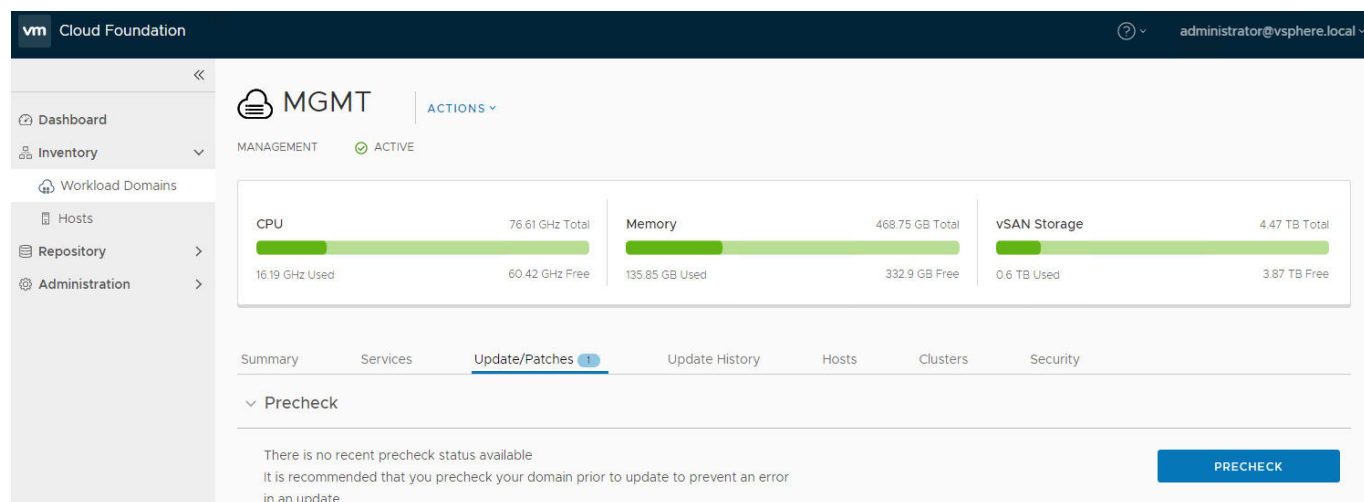
<input checked="" type="checkbox"/> Hardware compatibility - SCSI controller is VMware certified	
Description	Hardware compatibility - SCSI controller is VMware certified
Start Time	Sep 16, 2022, 10:44:11 AM
End Time	Sep 16, 2022, 10:44:12 AM
Health Status	 Silenced RESTORE PRECHECK

You can also silence failed vSAN prechecks in the SDDC Manager UI by clicking **Silence Precheck**. Silenced prechecks do not trigger warnings or block upgrades.

IMPORTANT

You should only silence alerts if you know that they are incorrect. Do not silence alerts for real issues that require remediation.


1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the workload domain where you want to run the precheck.
3. On the domain summary page, click the **Updates/Patches** tab. The image below is a sample screenshot and may not reflect the correct product versions.



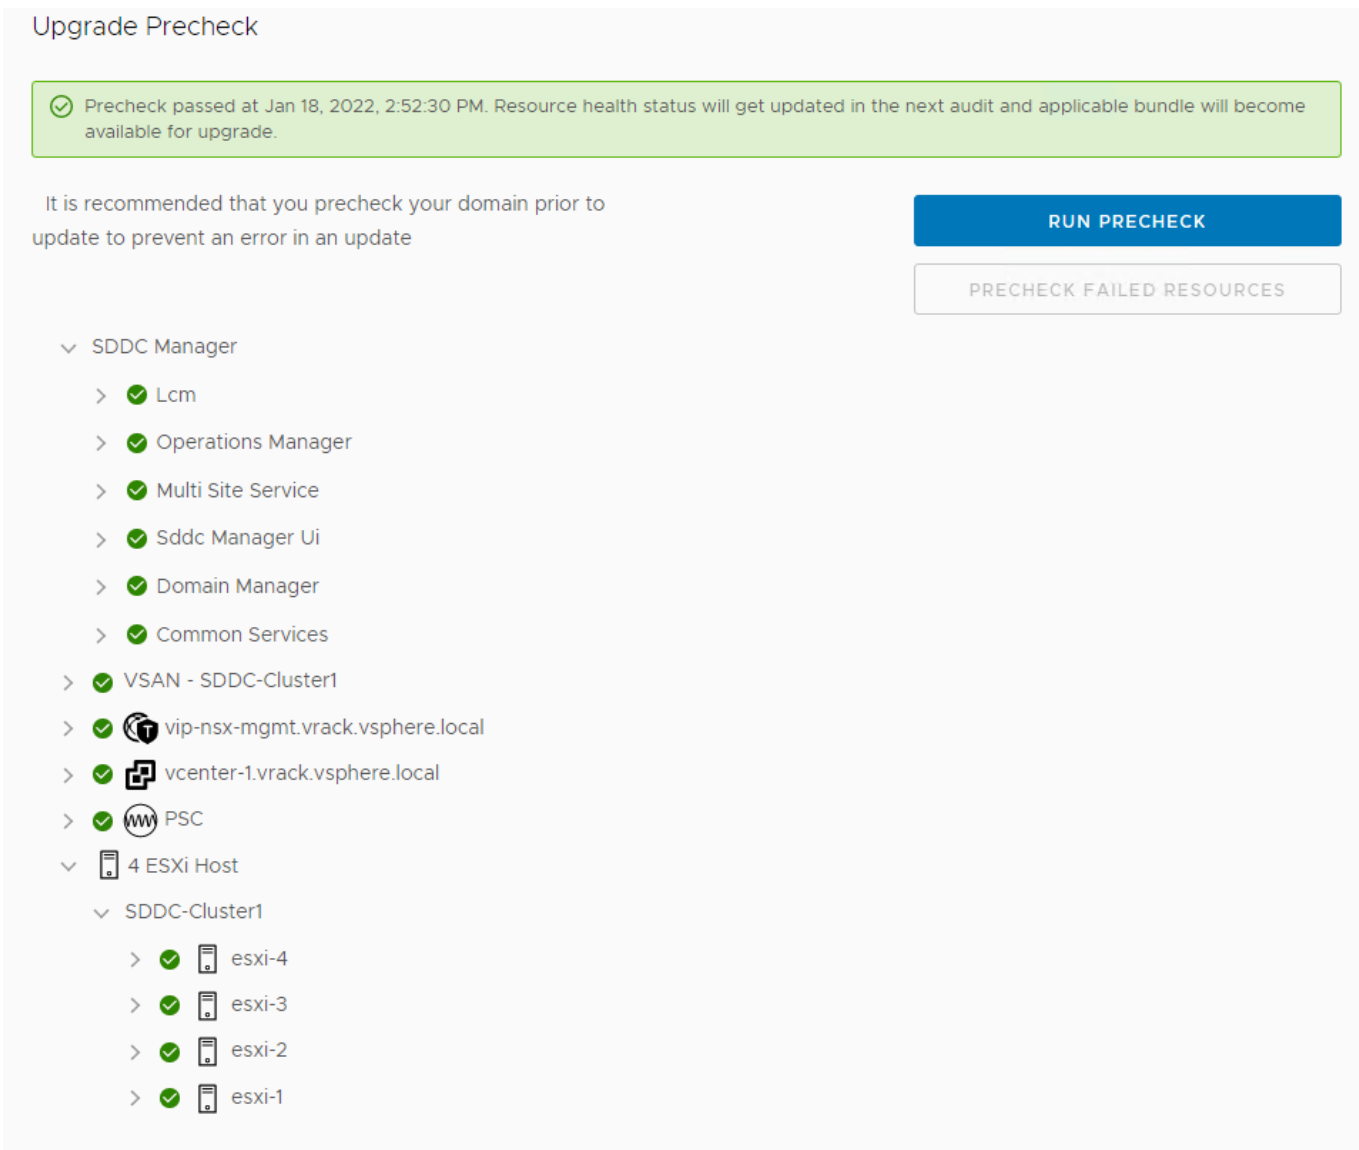
The screenshot shows the VMware Cloud Foundation SDDC Manager interface. The left navigation pane includes 'Dashboard', 'Inventory', 'Workload Domains', 'Hosts', 'Repository', and 'Administration'. The main content area shows the 'MGMT' management console with 'ACTIONS' and 'MANAGEMENT' tabs. Below this, there are three resource usage bars: CPU (76.61 GHz Total, 16.19 GHz Used, 60.42 GHz Free), Memory (468.75 GB Total, 135.85 GB Used, 332.9 GB Free), and vSAN Storage (4.47 TB Total, 0.6 TB Used, 3.87 TB Free). The 'Update/Patches' tab is selected, showing a 'Precheck' section with a message: 'There is no recent precheck status available. It is recommended that you precheck your domain prior to update to prevent an error in an update.' A blue 'PRECHECK' button is located at the bottom right of the precheck section.

4. Click **Precheck** to validate that the environment is ready to be upgraded.

Once the precheck begins, a message appears indicating the time at which the precheck was started.



5. Click **View Status** to see detailed tasks and their status. The image below is a sample screenshot and may not reflect the correct versions.



6. To see details for a task, click the Expand arrow.

If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **Precheck Failed Resources** to retry all failed tasks.

7. If ESXi hosts display a driver incompatibility issue when updating a VI workload domain using vSphere Lifecycle Manager baselines, perform the following steps:
 1. Identify the controller with the HCL issue.
 2. For the given controller, identify the supported driver and firmware versions on the source and target ESXi versions.
 3. Upgrade the firmware, if required.
 4. Upgrade the driver manually on the ESXi host and retry the task at which the upgrade failed.

8. If the workload domain contains a host that includes pinned VMs, the precheck fails at the Enter Maintenance Mode step. If the host can enter maintenance mode through vCenter Server UI, you can suppress this check for NSX and ESXi in VMware Cloud Foundation by following the steps below.
 1. Log in to SDDC Manager by using a Secure Shell (SSH) client with the user name `vcf` and password you specified in the deployment parameter workbook.
 2. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
 3. Add the following line to the end of the file:


```
lcm.nsxt.suppress.dry.run.emm.check=true
```

```
lcm.esx.suppress.dry.run.emm.check.failures=true
```
 4. Restart Lifecycle Management by typing the following command in the console window.


```
systemctl restart lcm
```
 5. After Lifecycle Management is restarted, run the precheck again.

The precheck result is displayed at the top of the Upgrade Precheck Details window. If you click **Exit Details**, the precheck result is displayed at the top of the Precheck section in the Updates/Patches tab.

Ensure that the precheck results are green before proceeding. A failed precheck may cause the update to fail.

Perform Update Precheck in SDDC Manager

You must perform a precheck in SDDC Manager before applying an update bundle to ensure that your environment is ready for the update.


Bundle-level pre-checks for vCenter are available in VMware Cloud Foundation.

NOTE

Because ESXi bundle-level pre-checks only work in minor-version upgrades (for example: from ESXi 7.x through 7.y, or from ESXi 8.x through 8.y), these prechecks do not run in VMware Cloud Foundation.

If you silence a vSAN Skyline Health alert in the vSphere Client, SDDC Manager skips the related precheck and indicates which precheck it skipped. Click **RESTORE PRECHECK** to include the silenced precheck. For example:

RESTORE PRECHECK

Resource Name	Description	Health Status	Error Description	Impact	Remediation
SDDC-Cluster1	Checks the age of the VMware Hardware Compatibility Guide database used for the HCL checks. Shows warning or error when it is older than 90 or 180 days, respectively. VMware updates the VCG frequently, so it is important to keep the local copy up-to-date	 SILENCED	Check skipped because 'com.vmware.vsan.health.test.hcldbuptodate' is silenced in vSAN		If you would like to run the check, click the Restore button or enable it through the vSphere UI

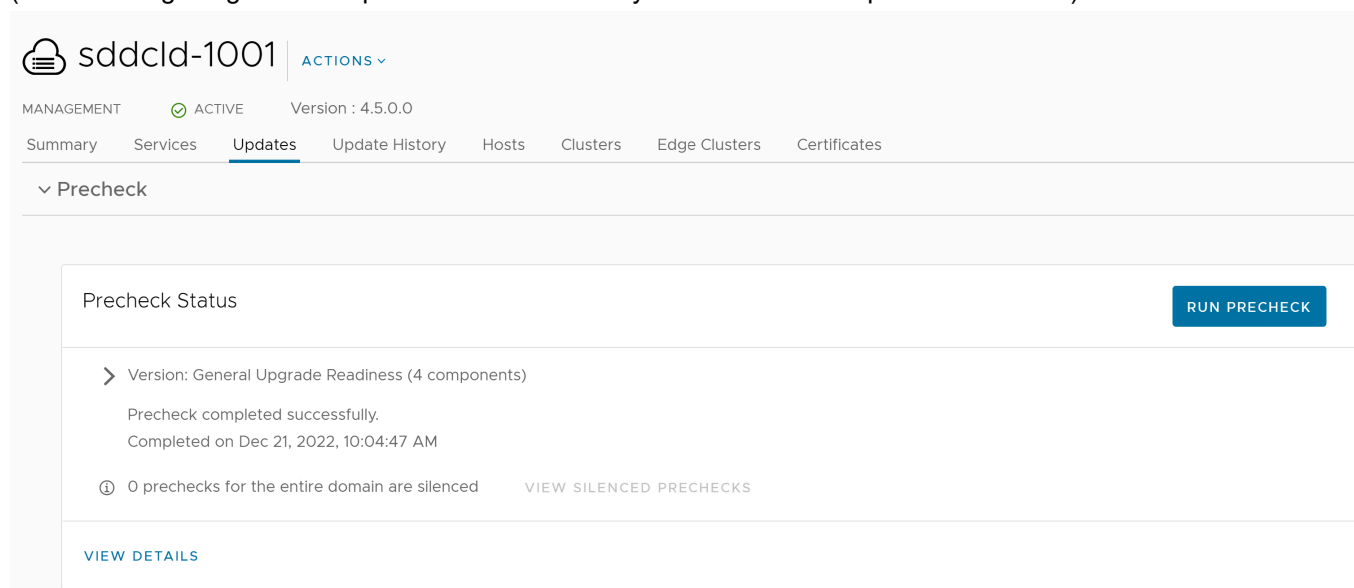
You can also silence failed vSAN prechecks in the SDDC Manager UI by clicking **Silence Precheck**. Silenced prechecks do not trigger warnings or block upgrades.

IMPORTANT

Only silence alerts if you know that they are incorrect. Do not silence alerts for real issues that require remediation.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the workload domain where you want to run the precheck.
3. On the domain summary page, click the **Updates** tab.

(The following image is a sample screenshot and may not reflect current product versions.)



The screenshot shows the SDDC Manager interface for workload domain 'sddcId-1001'. The 'Updates' tab is active, and the 'Precheck' section is expanded. The 'Precheck Status' section shows a successful precheck completion on Dec 21, 2022, at 10:04:47 AM. A 'RUN PRECHECK' button is located in the top right corner of the Precheck Status section. Below the status, there is a message indicating that 0 prechecks for the entire domain are silenced, with a link to 'VIEW SILENCED PRECHECKS'. A 'VIEW DETAILS' link is also present at the bottom of the Precheck Status section.

NOTE

It is recommended that you Precheck your workload domain prior to performing an upgrade.

4. Click **RUN PRECHECK** to select the components in the workload domain you want to precheck.
 - a) You can select to run a Precheck only on vCenter or the vSphere cluster. All components in the workload domain are selected by default. To perform a precheck on certain components, choose **Custom selection**.

Precheck

Run a general precheck or precheck for a specific release version. You can precheck across the entire workload domain or select specific components.

Target Version ⓘ General Upgrade Readiness ▾

Run Precheck on Entire Workload Domain Custom selection

Datagrid lists the components that are going to be prechecked.

<input type="checkbox"/>	Component	Description
<input checked="" type="checkbox"/>	sddc-manager.vrack.vsphere.local	SDDC_MANAGER
<input checked="" type="checkbox"/>	vip-nsx-mgmt.vrack.vsphere.local	NSX
<input checked="" type="checkbox"/>	SDDC-Cluster1	CLUSTER
<input checked="" type="checkbox"/>	vcenter-1.vrack.vsphere.local	VCENTER
<input checked="" type="checkbox"/>	4	

4 objects

RUN PRECHECK
CANCEL

- b) If there are pending upgrade bundles available, then the "Target Version" dropdown contains "General Upgrade Readiness" and the available VMware Cloud Foundation versions to upgrade to. If there is an available VMware Cloud Foundation upgrade version, there will be extra checks - bundle-level prechecks for hosts, vCenter Server, and so forth. The version specific prechecks will only run prechecks on components that have available upgrade bundles downloaded.

sddcId-1001 | ACTIONS ▾

MANAGEMENT ACTIVE Version : 4.5.0.0

Starting VCF 4.4 release, SSH Service on ESXI hosts is disabled on new and upgraded VCF deployments. ✕

Summary Services **Updates** Update History Hosts Clusters Edge Clusters Certificates

[< BACK TO UPDATES](#)

Precheck

Run a general precheck or precheck for a specific release version. You can precheck across the entire workload domain or select specific components.

Target Version ⓘ 5.0.0.0 ▾

Precheck Scope ⓘ All components in version - 5.0.0.0 ▾

All components in version - 5.0.0.0
 SDDC_MANAGER_VCF - 5.0.0.0-20774436
 SDDC_MANAGER_VCF - 5.0.0.0-20774436
 NSX - 4.0.2.0.0-20606912
VCENTER - 8.0.0.20000-20646705
 ESX_HOST1 - 8.0.0-20646706
 Custom selection

to be prechecked. The order of the items in the datagrid is following the upgrade sequence in the selected version.

<input type="checkbox"/>	Component	Description	Target Version
<input checked="" type="checkbox"/>	vip-nsx-mgmt.vrack.vsphere.local	NSX	4.0.2.0.0-20606912
<input checked="" type="checkbox"/>	vcenter-1.vrack.vsphere.local	VCENTER	8.0.0.20000-20646705
<input checked="" type="checkbox"/>	SDDC-Cluster1	CLUSTER	8.0.0-20646706
<input checked="" type="checkbox"/>	4		


4 objects

RUN PRECHECK
CANCEL

5. When the precheck begins, a progress message appears indicating the precheck progress and the time when the precheck began.

Precheck

In Progress

 72% Completed

Started on Dec 21, 2022, 10:02:29 AM

Precheck report will be shown once completed.





NOTE

Parallel precheck workflows are supported. If you want to precheck multiple domains, you can repeat steps 1-5 for each of them without waiting for step 5 to finish.

6. Once the Precheck is complete, the report appears. Click through **ALL**, **ERRORS**, **WARNINGS**, and **SILENCED** to filter and browse through the results.

Precheck

Results Completed on Dec 21, 2022, 10:04:47 AM





 174 Passed  7 Errors  3 Warnings  0 Silenced


[RETRY ALL FAILED RESOURCES](#)

Report

ALL ERRORS WARNINGS SILENCED

Version: General Upgrade Readiness (4 components)

- >  SDDC Manager sddc-manager.vrack.vsphere.local
- >  NSX vip-nsx-mgmt.vrack.vsphere.local
- >  ESXi Host Cluster SDDC-Cluster1
- >  vCenter vcenter-1.vrack.vsphere.local

Resource Name	Description	Health Status	Error Description	Impact	Remediation
 Precheck entry not selected					

7. To see details for a task, click the expander arrow.
- If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **RETRY ALL FAILED RESOURCES** to retry all failed tasks.
8. If ESXi hosts display a driver incompatibility issue when updating a VI workload domain using vSphere Lifecycle Manager baselines, perform the following steps:
1. Identify the controller with the HCL issue.
 2. For the given controller, identify the supported driver and firmware versions on the source and target ESXi versions.
 3. Upgrade the firmware, if required.
 4. Upgrade the driver manually on the ESXi host and retry the task at which the upgrade failed.
9. If the workload domain contains a host that includes pinned VMs, the precheck fails at the Enter Maintenance Mode step. If the host can enter maintenance mode through vCenter Server UI, you can suppress this check for NSX and ESXi in VMware Cloud Foundation by following the steps below.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client with the user name `vcf` and password.
2. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
3. Add the following line to the end of the file:

```
lcm.nsxt.suppress.dry.run.emm.check=true
```

```
lcm.esx.suppress.dry.run.emm.check.failures=true
```

4. Restart Lifecycle Management by typing the following command in the console window.

```
systemctl restart lcm
```

5. After Lifecycle Management is restarted, run the precheck again.

The precheck result is displayed at the top of the Upgrade Precheck Details window. If you click **Exit Details**, the precheck result is displayed at the top of the Precheck section in the Updates tab.

Ensure that the precheck results are green before proceeding. Although a failed precheck will not prevent the upgrade from proceeding, it may cause the update to fail.

Apply the VMware Cloud Foundation 5.2.x Upgrade Bundle

The VMware Cloud Foundation Upgrade bundle upgrades the SDDC Manager appliance and VMware Cloud Foundation services.

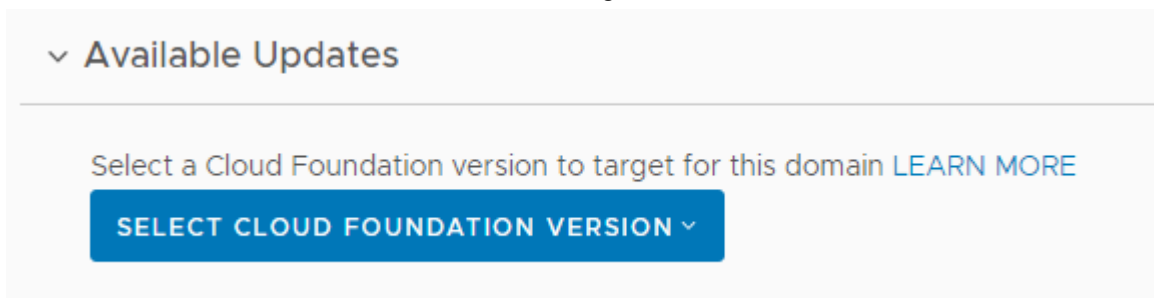
- Download the VMware Cloud Foundation update bundle for your target release. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Ensure you have a recent successful backup of SDDC Manager using an external SFTP server.
- Ensure you have taken a snapshot of the SDDC Manager appliance.
- Ensure you have recent successful backups of the components managed by SDDC Manager.
- [Perform Update Precheck in SDDC Manager](#) and resolve any issues.

After SDDC Manager is upgraded to 5.2 or later, new functionality is introduced that allows you to upgrade SDDC Manager without having to upgrade the entire VMware Cloud Foundation BOM. See [Independent SDDC Manager Upgrade using the SDDC Manager UI](#).

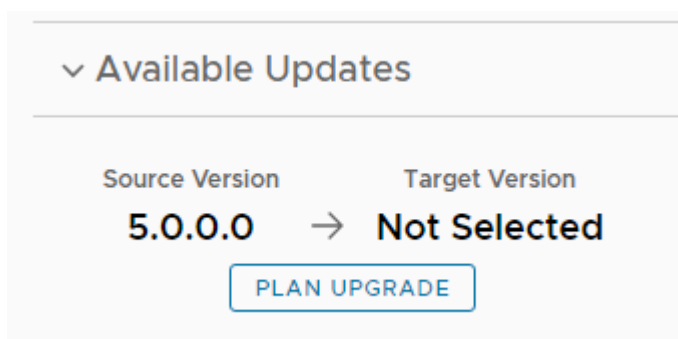
1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. On the Workload Domains page, click the management domain and then click the **Updates** tab.
3. In the **Available Updates** section, select the target VMware Cloud Foundation release or click **Plan Upgrade**.

The available options depend on the source version of VMware Cloud Foundation.

- For VMware Cloud Foundation 4.5.x, select the target version.



- For VMware Cloud Foundation 5.x, click **Plan Upgrade**, select a target version, and click **Confirm**.



4. Click **Update Now** or **Schedule Update** next to the VMware Cloud Foundation Upgrade bundle.
5. If you selected **Schedule Update**, select the date and time for the bundle to be applied and click **Schedule**.

Schedule Update ✕

The bundle will be scheduled based on your selected date and time.

Date

Time

CANCEL

SCHEDULE

If you clicked **Update Now**, the VMware Cloud Foundation Update Status window displays the components that will be upgraded and the upgrade status. Click **View Update Activity** to view the detailed tasks. After the upgrade is completed, a green bar with a check mark is displayed.

6. Click **Finish**.
When the update completes successfully, you are logged out of the SDDC Manager UI and must log in again.

Apply VMware Cloud Foundation Configuration Updates

VMware Cloud Foundation Configuration Updates identifies and resolves any discrepancies between the intended/prescribed configuration and the actual configuration, ensuring that the deployment aligns with the recommended configuration. This process includes reconciling the configuration for 2nd party software components listed in the VMware Cloud Foundation Bill of Materials (BOM).

Configuration updates may be required after you apply software updates. Once a configuration update becomes available, you can apply it immediately or wait until after you have applied all software updates. Configuration Updates must be performed during a maintenance window.

Configuration Updates can be applied to multiple domains in parallel. However, if a Configuration Update is in progress, another configuration update on the same domain should not be attempted.

NOTE

Configuration Updates in VCF detects and reconciles to a prescribed configuration for the release. Once reconciled, it does not identify subsequent non-compliance arising from out of band changes.

The following configuration updates may become available, depending on your source version of VMware Cloud Foundation:

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
ConfigureVsanHaloIsolationAddressesConfigDrift	Configures the vSAN HA network isolation address to use the vSAN vmkernel interface gateway, in conformance with VCF best practices.	4.3.0.0	CLUSTER	FIX	vCenter 7.0.3
ToggleVSanRecommendationConfigDrift	Disables vSAN baseline recommendations for vSAN enabled clusters.	4.4.1.0	CLUSTER	FIX	vCenter 7.0.0
RemoveNfsDatastoreConfigDrift	Removes NFS datastore on hosts.	5.0.0.0	CLUSTER	FIX	NA
CloudAdminRoleConfigDrift	Creates Cloud Admin role in vCenter Server for the management domain.	5.0.0.0	DOMAIN	FEATURE	vCenter 7.0.3
AllowBrokerConfigurationConfigDrift	Adds <code>config.SDDC.Deployed.AllowBrokerConfiguration</code> advanced property in vCenter Server. This property restricts the user from configuring an external IDP from the vCenter UI in the ELM ring (workload domain vCenters). Configuration is only possible from the management domain vCenter UI and isolated workload domain vCenter UI.	5.1.0.0	DOMAIN	FEATURE	vCenter 8.0.2
ClusterHaSettingsConfigDrift	Removes <code>das.includeFTcomplianceChecks</code> option HA	5.1.0.0	DOMAIN	FEATURE	vCenter 8.0.1

Table continued on next page

Continued from previous page

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
	configuration from all clusters on the management domain.				
ComputeManager SettingsDrift	Creates an internal NSX service account to enable NSX to vSphere Lifecycle Manager communication.	5.1.0.0	DOMAIN	FEATURE	vCenter 7.0.2.00400, NSX 3.1.3.0.0
DvpgConfiguration Drift	Creates a new distributed virtual port group named VM_MANAGEMENT in the target domain, and migrates all VMs connected to the management port group to this new port group. The purpose of this feature is to allow separation of traffic coming from management VMs and ESXi hosts. VMs migrated: VCSA, SDDC Manager, NSX Manager and Edge VMs.	5.1.0.0	CLUSTER	FEATURE	NA
EsxAdvancedOptionsConfigDrift	Configures <code>UserVars.SuppressShellWarning</code> property on every ESXi host to false, to enable warnings for ESXi Shell and SSH services.	5.1.0.0	DOMAIN	FEATURE	NA
WorkspaceOneBrokerConfigDrift	Configures BOM components as OIDC relying	5.1.0.0	DOMAIN	FEATURE	vCenter 8.0.2, NSX 4.1.2

Table continued on next page

Continued from previous page

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
	parties of Workspace ONE Broker in vCenter.				
RegisterSDDCmanagerAsVCExtensionConfigDrift	Register SDDC Manager as an extension in a workload domain vCenter.	5.2.0.0	DOMAIN	FEATURE	vCenter 7.0.0

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the workload domain name and then click the **Updates** tab.
3. Click **Run Precheck** to run the upgrade precheck.
Resolve any issues before proceeding with the upgrade.
4. Expand **Available Configuration Updates**, click **Apply All**.

▼ Available Configuration Updates +

Configuration updates may be required after you apply software updates. Once a configuration update becomes available, you can apply it immediately or wait until after you have applied all software updates.

APPLY ALL

Description	Type	Resource Type
Remove das.includeFTcomplianceChecks option from all clusters HA configuration	FEATURE	DOMAIN
ESXi advanced options for ESXi Hosts Addition Drift	FEATURE	DOMAIN
Register SDDC Manager as an extension in a domain vCenter	FEATURE	DOMAIN
Creates a Distributed Virtual Port Group to enable traffic isolation between management VMs and ESXi hosts	FEATURE	CLUSTER

- **FEATURE:** Configuration change required for a new feature.
- **FIX:** Configuration change associated with a fix for a defect.

5. Check the progress of a configuration update by clicking the task in the Tasks panel.

Tasks ⊞ ↻ ×

REFRESH RESET FILTERS

Task	Subtask	Task Status	Last Occurrence
Configuration Updates	DVPG Configuration Drift	<div style="width: 50%; height: 10px; background-color: #0070C0; margin-bottom: 2px;"></div> 50%	5/19/23, 10:55 AM

6. After the configuration updates are successfully applied, they will no longer appear in the table.

▼ Available Configuration Updates

✔ The recent configuration update successfully completed, please check upgrade details. View Details

No Available Updates.

Pending Configuration Updates do not block future BOM upgrades.

Upgrade VMware Aria Suite Lifecycle and VMware Aria Suite Products for VMware Cloud Foundation

VMware Cloud Foundation does not manage upgrades for VMware Aria Suite Lifecycle and the VMware Aria Suite products. Use VMware Aria Suite Lifecycle to upgrade VMware Aria Suite products.

If you had VMware Aria Suite Lifecycle, VMware Aria Operations for Logs, VMware Aria Automation, VMware Aria Operations, or Workspace ONE Access in your pre-upgrade environment, you must upgrade them from VMware Aria Suite Lifecycle.

Use VMware Aria Suite Lifecycle to:

- Download upgrade binaries
- Create snapshots of the virtual appliances
- Run pre-upgrade checks
- Upgrade VMware Aria Suite products

You can upgrade VMware Aria Suite products as new versions become available in VMware Aria Suite Lifecycle. VMware Aria Suite Lifecycle will only allow upgrades to compatible and supported versions of VMware Aria Suite products.

NOTE

See the [VMware Interoperability Matrix](#) for information about which versions are supported with your version of VMware Cloud Foundation and [KB 88829](#) for more information about supported upgrade paths using VMware Aria Suite Lifecycle.

IMPORTANT

The VMware Cloud Foundation 5.2 BOM requires VMware Aria Suite Lifecycle 8.18 or higher.

NOTE

The VMware Aria Suite of products were formerly known as the vRealize Suite of products.

1. Log in to VMware Aria Suite Lifecycle at `https://<aria_suite_lifecycle_manager_fqdn>` as the administrator.
2. Upgrade VMware Aria Suite products.

Upgrade VMware Aria Suite Lifecycle first and then upgrade VMware Aria Suite products.

See “Upgrading VMware Aria Suite Lifecycle and VMware Aria Suite Products” in the *VMware Aria Suite Lifecycle Installation, Upgrade, and Management Guide* for your current version of [VMware Aria Suite Lifecycle](#).

Upgrade NSX for VMware Cloud Foundation in a Federated Environment

If NSX Federation is configured between two VMware Cloud Foundation instances, SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must manually upgrade the NSX Global Managers for each instance.

Download NSX Global Manager Upgrade Bundle

SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must download the NSX upgrade bundle manually to upgrade the NSX Global Managers.

1. Log in to the Broadcom Support Portal and browse to **My Downloads > VMware NSX**.
2. Click the version of NSX to which you are upgrading.
3. Locate the **NSX *version* Upgrade Bundle** and verify that the upgrade bundle filename extension ends with `.mub`.
The upgrade bundle filename has the following format `VMware-NSX-upgrade-bundle-versionnumber.buildnumber.mub`.
4. Click the download icon to download the upgrade bundle to the system where you access the NSX Global Manager UI.

Upgrade the Upgrade Coordinator for NSX Federation

The upgrade coordinator runs in the NSX Manager. It is a self-contained web application that orchestrates the upgrade process of hosts, NSX Edge cluster, NSX Controller cluster, and the management plane.

The upgrade coordinator guides you through the upgrade sequence. You can track the upgrade process and, if necessary, you can pause and resume the upgrade process from the UI.

1. In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Proceed to Upgrade**.
4. Navigate to the upgrade bundle `.mub` file you downloaded or paste the download URL link.
 - Click **Browse** to navigate to the location you downloaded the upgrade bundle file.
 - Paste the VMware download portal URL where the upgrade bundle `.mub` file is located.
5. Click **Upload**.
When the file is uploaded, the **Begin Upgrade** button appears.
6. Click **Begin Upgrade** to upgrade the upgrade coordinator.

NOTE

Upgrade one upgrade coordinator at a time.

7. Read and accept the EULA terms and accept the notification to upgrade the upgrade coordinator..
8. Click **Run Pre-Checks** to verify that all NSX components are ready for upgrade.
The pre-check checks for component connectivity, version compatibility, and component status.
9. Resolve any warning notifications to avoid problems during the upgrade.

Upgrade NSX Global Managers for VMware Cloud Foundation

Manually upgrade the NSX Global Managers when NSX Federation is configured between two VMware Cloud Foundation instances.

Before you can upgrade NSX Global Managers, you must upgrade all VMware Cloud Foundation instances in the NSX Federation, including NSX Local Managers, using SDDC Manager.

1. In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Start** to upgrade the management plane and then click **Accept**.
4. On the Select Upgrade Plan page, select **Plan Your Upgrade** and click **Next**.

The NSX Manager UI, API, and CLI are not accessible until the upgrade finishes and the management plane is restarted.

Upgrade NSX for VMware Cloud Foundation 5.2.x

Upgrade NSX in the management domain and VI workload domains. VMware Cloud Foundation 5.2.1 supports in-place host upgrades for clusters that use vSphere Lifecycle Manager baselines.

Until SDDC Manager is upgraded to version 5.2, you must upgrade NSX in the management domain before you upgrade NSX in a VI workload domain. Once SDDC Manager is at version 5.2 or later, you can upgrade NSX in VI workload domains before or after upgrading NSX in the management domain.

Upgrading NSX involves the following components:

- Upgrade Coordinator
- NSX Edges/Clusters (if deployed)
- Host clusters
- NSX Manager cluster

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates/Patches** tab.

When you upgrade NSX components for a selected VI workload domain, those components are upgraded for all VI workload domains that share the NSX Manager cluster.

3. Click **Precheck** to run the upgrade precheck.
Resolve any issues before proceeding with the upgrade.

NOTE

The NSX precheck runs on all VI workload domains in your environment that share the NSX Manager cluster.

4. For VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for NSX.

Available Updates 3

Target Version	Progress	ACTIONS
5.2.0.0	1 of 4 steps done	VIEW DETAILS
VMware Software Update 5.2.0.0 Released Jun 19, 2024 9 GB The upgrade bundle for VMware NSX Data Center 4.2.0.0. Customers are strongly encouraged to run the NSX Upgrade Evaluation Tool. For more information, see https://docs.vmware.com/en/VMware-NSX/4.2/rn/vmware-nsxt-data-center-42-release-notes/index.html .		SCHEDULE UPDATE UPDATE NOW

[View Details](#)

- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.
By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.
- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.
By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) If you selected the **Schedule Upgrade** option, specify the date and time for the NSX bundle to be applied and click **Next**.
- f) On the Review page, review your settings and click **Finish**.

If you selected **Upgrade Now**, the NSX upgrade begins and the upgrade components are displayed. The upgrade view displayed here pertains to the workload domain where you applied the bundle. Click the link to the associated workload domains to see the components pertaining to those workload domains. If you selected **Schedule Upgrade**, the upgrade begins at the time and date you specified.

5. For VMware Cloud Foundation 5.2.1:

- a) In the Available Updates section, click the **Configure Update** button.

- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.

By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.

- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.

By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.

By default ESXi hosts are placed into maintenance mode during an upgrade. Starting with VMware Cloud Foundation 5.2.1, in-place upgrades are available for workload domains in which all the clusters use vSphere Lifecycle Manager baselines. If NSX Manager is shared between workload domains, in-place upgrade is only available if all the clusters in all the workload domains that share the NSX Manager use vLCM baselines. If the option is available, you can select **In-place** as the upgrade mode to avoid powering off and placing hosts into maintenance mode before the upgrade.

NOTE

To perform an in-place upgrade, the target NSX version must be the VMware Cloud Foundation 5.2.1 BOM version or later.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) On the Review page, review your settings and click **Run Precheck**.

The precheck begins. Resolve any issues until the precheck succeeds.

- f) After the precheck succeeds, click **Schedule Update** and select an option.
6. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).

If a component upgrade fails, the failure is displayed across all associated workload domains. Resolve the issue and retry the failed task.

When all NSX workload components are upgraded successfully, a message with a green background and check mark is displayed.

Upgrade vCenter Server for VMware Cloud Foundation 5.2.x

The upgrade bundle for VMware vCenter Server is used to upgrade the vCenter Server instances managed by SDDC Manager. Upgrade vCenter Server in the management domain before upgrading vCenter Server in VI workload domains.

- Download the VMware vCenter Server upgrade bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Take a file-based backup of the vCenter Server appliance before starting the upgrade. See [Manually Back Up vCenter Server](#).

NOTE

After taking a backup, do not make any changes to the vCenter Server inventory or settings until the upgrade completes successfully.

- If your workload domain contains Workload Management (vSphere with Tanzu) enabled clusters, the supported target release depends on the version of Kubernetes (K8s) currently running in the cluster. Older versions of K8s might require a specific upgrade sequence. See [KB 92227](#) for more information.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

4. Upgrading to VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for vCenter Server.
 - b) Click **Confirm** to confirm that you have taken a file-based backup of the vCenter Server appliance before starting the upgrade.
 - c) If you selected **Schedule Update**, click the date and time for the bundle to be applied and click **Schedule**.
 - d) If you are upgrading from VMware Cloud Foundation 4.5.x, enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.

Schedule Update

- 1 Introduction
- 2 Temporary Network
- 3 Review

Temporary Network ×

The migration-based upgrade requires a temporary address to be used for a parallel instance. This will only be used during the upgrade.

IP Address

Subnet Mask

Gateway

- e) Review the upgrade settings and click **Finish**.
5. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 5.x:
- a) In the Available Updates section, click **Configure Update**.
 - b) Select the upgrade mechanism and click **Next**.

Option	Description
vCenter Reduced Downtime Upgrade	<p>The reduced downtime upgrade process uses a migration-based approach. In this approach, a new vCenter Server Appliance is deployed and the current vCenter data and configuration is copied to it.</p> <p>During the preparation phase of a reduced downtime upgrade, the source vCenter Server Appliance and all resources remain online. The only downtime occurs when the source vCenter Server Appliance is stopped, the configuration is switched over to the target vCenter, and the services are started. The downtime is expected to take approximately 5 minutes under ideal network, CPU, memory, and storage provisioning.</p> <p style="text-align: center;">NOTE To perform a vCenter Reduced Downtime Upgrade, the target vCenter version must be the VMware Cloud Foundation 5.2.1 BOM version or later.</p>

Table continued on next page

Continued from previous page

Option	Description
vCenter Regular Upgrade	During a regular upgrade, the vCenter Server Appliance is offline for the duration of the upgrade.

- c) Select a backup option and click **Next**.
- d) For an RDU update, provide a temporary network to be used only during the upgrade and click **Next**.

Automatic	Automatically assign network information.
Static	Enter an IP address, subnet mask, and gateway. The IP address must be in the management subnet.

- e) Schedule the update and click **Next**.

For vCenter Reduced Downtime Upgrade	Select scheduling options for the preparation and switchover phases of the upgrade. NOTE If you are scheduling the switchover phase, you must allow a minimum of 4 hours between the start of preparation and the start of switchover.
For vCenter Regular Upgrade	Select an Upgrade Now or Schedule Update .

- f) Review the upgrade settings and click **Finish**.
6. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 4.5.x:
 - a) In the Available Updates section, click **Configure Update**.
 - b) Enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.
 - c) Select a backup option and click **Next**.
 - d) Schedule the update and click **Next**.
 - e) Review the upgrade settings and click **Finish**.
 7. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).
 8. After the upgrade is complete, remove the old vCenter Server appliance (if applicable).

NOTE

Removing the old vCenter is only required for major upgrades. If you performed a vCenter RDU patch upgrade, the old vCenter is automatically removed after a successful upgrade.

If the upgrade fails, resolve the issue and retry the failed task. If you cannot resolve the issue, restore vCenter Server using the file-based backup. See [Restore vCenter Server](#). vCenter RDU upgrades perform automatic rollback if the upgrade fails.

Once the upgrade successfully completes, use the vSphere Client to change the vSphere DRS Automation Level setting back to the original value (before you took a file-based backup) for each vSphere cluster that is managed by the vCenter Server. See [KB 87631](#) for information about using VMware PowerCLI to change the vSphere DRS Automation Level.

Upgrade ESXi for VMware Cloud Foundation 5.2.1

VMware Cloud Foundation 5.2.1 and later support workload domains that include vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters. There is a single procedure for upgrading both vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters.

- Validate that the ESXi passwords are valid.
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.
- For clusters that use vSphere Lifecycle Manager images:
 - You must upgrade NSX and vCenter Server before you can upgrade ESXi hosts with a vSphere Lifecycle Manager image.
 - If you want to add firmware to the vSphere Lifecycle Manager image, you must install the Hardware Support Manager from your vendor. See [Firmware Updates](#).
 - A supported vSphere Lifecycle Manager image must be available in SDDC Manager. See steps 1-3 in [Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2](#) for more information.
- For clusters that use vSphere Lifecycle Manager baselines, download the ESXi bundle. See [Downloading Upgrade Bundles](#).

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager images when your target version is VMware Cloud Foundation 5.2, see [Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2](#).

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager baselines when your target version is VMware Cloud Foundation 5.2, see [Upgrade ESXi with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation 5.2](#).

If you are using external (non-vSAN) storage, the following procedure updates the ESXi hosts attached to the external storage. However, updating and patching the storage software and drivers is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability after an ESXi upgrade, consult the vSphere HCL and your storage vendor.

For clusters that use vSphere Lifecycle Manager baselines:

- If you want to skip any hosts while applying an ESXi update a workload domain, you must add these hosts to the `application-prod.properties` file before you begin the update. See "Skip Hosts During ESXi Update".
- To perform ESXi upgrades with custom ISO images or async drivers see "Upgrade ESXi with Custom ISOs" and "Upgrade ESXi with Stock ISO and Async Drivers".

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

4. In the Available Updates section, click **Configure Update**.
5. Read the introductory information and click **Next**.
6. If any clusters in the workload domain use vSphere Lifecycle Manager images, select the clusters to update and click **Next**.
7. Assign an image to each cluster that uses vSphere Lifecycle Manager images and click **Next**.
8. If any clusters in the workload domain use vSphere Lifecycle Manager baselines, select the clusters to upgrade and click **Next**.

The default setting is to upgrade all clusters. To upgrade specific clusters, select **Custom selection** and select the clusters to upgrade.

9. If the workload domain you are upgrading only includes clusters that use vSphere Lifecycle Manager baselines, select a scheduling option.
10. Select the upgrade options and click **Next**.

By default, the selected clusters are upgraded in parallel. If you selected more than ten clusters to be upgraded, the first ten are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Select **Enable Quick Boot** to reduce the upgrade time by skipping the physical reboot of the host.

Select **Migrate Powered Off and Suspended VMs** to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster.

For clusters that use vSphere Lifecycle Manager images, select **Enforce Live Patch** when the cluster image includes a Live Patch. With the **Enforce Live Patch** option, vSphere Lifecycle Manager does not place the hosts in the cluster into maintenance mode, hosts are not rebooted, and there is no need to migrate the virtual machines running on the hosts in the cluster.

11. Review the settings, and click **Finish** or **Run Precheck**.

If the upgrade includes any clusters that use vSphere Lifecycle Manager images VMware Cloud Foundation runs a cluster image hardware compatibility and compliance precheck. Resolve any reported issues before proceeding.

12. After the precheck succeeds, click **Schedule Update**, select a scheduling option, and click **Finish**.

13. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [Upgrade vSAN on-disk format versions](#).

Upgrade ESXi with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation 5.2

The following procedure describes upgrading ESXi hosts in workload domains that use vSphere Lifecycle Manager baselines when your target version is VMware Cloud Foundation 5.2.

- Validate that the ESXi passwords are valid.
- Download the ESXi bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager images when your target version is VMware Cloud Foundation 5.2, see [Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2](#).

VMware Cloud Foundation 5.2.1 and later support workload domains that include vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters. If you are upgrading to VMware Cloud Foundation 5.2.1, see [Upgrade ESXi for VMware Cloud Foundation 5.2.1](#).

By default, the upgrade process upgrades the ESXi hosts in all clusters in a workload domain in parallel. If you have multiple clusters in a workload domain, you can select the clusters to upgrade.

If you want to skip any hosts while applying an ESXi update a workload domain, you must add these hosts to the `application-prod.properties` file before you begin the update. See "Skip Hosts During ESXi Update".

To perform ESXi upgrades with custom ISO images or async drivers see "Upgrade ESXi with Custom ISOs" and "Upgrade ESXi with Stock ISO and Async Drivers".

If you are using external (non-vSAN) storage, the following procedure updates the ESXi hosts attached to the external storage. However, updating and patching the storage software and drivers is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability after an ESXi upgrade, consult the vSphere HCL and your storage vendor.

1. Navigate to the **Updates/Patches** tab of the workload domain.
2. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

3. In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for ESXi.

If you selected **Schedule Update**, click the date and time for the bundle to be applied and click **Schedule**.

4. Select the clusters to upgrade and click **Next**.

The default setting is to upgrade all clusters. To upgrade specific clusters, click **Enable cluster-level selection** and select the clusters to upgrade.

5. Click **Next**.

6. Select the upgrade options and click **Next**.

By default, the selected clusters are upgraded in parallel. If you selected more than ten clusters to be upgraded, the first ten are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Select **Enable Quick Boot** to reduce the upgrade time by skipping the physical reboot of the host.

Select **Migrate Powered Off and Suspended VMs** to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster.

7. On the Review page, review your settings and click **Finish**.
8. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [Upgrade vSAN on-disk format versions](#).

Upgrade vSAN Witness Host for VMware Cloud Foundation

If your VMware Cloud Foundation environment contains stretched clusters, update and remediate the vSAN witness host.

Download the ESXi ISO that matches the version listed in the the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes*.

1. In a web browser, log in to vCenter Server at https://vcenter_server_fqdn/ui.
2. Upload the ESXi ISO image file to vSphere Lifecycle Manager.
 - a) Click **Menu > Lifecycle Manager**.
 - b) Click the **Imported ISOs** tab.
 - c) Click **Import ISO** and then click **Browse**.
 - d) Navigate to the ESXi ISO file you downloaded and click **Open**.
 - e) After the file is imported, click **Close**.
3. Create a baseline for the ESXi image.
 - a) On the Imported ISOs tab, select the ISO file that you imported, and click **New baseline**.
 - b) Enter a name for the baseline and specify the **Content Type** as Upgrade.
 - c) Click **Next**.
 - d) Select the ISO file you had imported and click **Next**.
 - e) Review the details and click **Finish**.
4. Attach the baseline to the vSAN witness host.
 - a) Click **Menu > Hosts and Clusters**.
 - b) In the Inventory panel, click **vCenter > Datacenter**.

- c) Select the vSAN witness host and click the **Updates** tab.
- d) Under Attached Baselines, click **Attach** › **Attach Baseline or Baseline Group**.
- e) Select the baseline that you had created in step 3 and click **Attach**.
- f) Click **Check Compliance**.

After the compliance check is completed, the **Status** column for the baseline is displayed as Non-Compliant.

5. Remediate the vSAN witness host and update the ESXi hosts that it contains.
 - a) Right-click the vSAN witness and click **Maintenance Mode** › **Enter Maintenance Mode**.
 - b) Click **OK**.
 - c) Click the **Updates** tab.
 - d) Select the baseline that you had created in step 3 and click **Remediate**.
 - e) In the End user license agreement dialog box, select the check box and click **OK**.
 - f) In the Remediate dialog box, select the vSAN witness host, and click **Remediate**.

The remediation process might take several minutes. After the remediation is completed, the **Status** column for the baseline is displayed as Compliant.

- g) Right-click the vSAN witness host and click **Maintenance Mode** › **Exit Maintenance Mode**.
- h) Click **OK**.

Skip Hosts During ESXi Update

You can skip hosts while applying an ESXi update to a workload domain. The skipped hosts are not updated.

NOTE

You cannot skip hosts that are part of a VI workload domain that is using vSphere Lifecycle Manager images, since these hosts are updated at the cluster-level and not the host-level.

1. Using SSH, log in to the SDDC Manager appliance with the user name `vcmf` and password you specified in the deployment parameter sheet.
2. Type `su` to switch to the root account.
3. Retrieve the host IDs for the hosts you want to skip.

```
curl 'https://SDDC_MANAGER_IP/v1/hosts' -i -u 'username:password' -X GET -H 'Accept: application/json' |json_pp
```

Replace the SDDC Manager FQDN, user name, and password with the information for your environment.

4. Copy the ids for the hosts you want to skip from the output. For example:

```
...
    "fqdn" : "esxi-2.vrack.vsphere.local",
    "esxiVersion" : "6.7.0-16075168",
    "id" : "b318fe37-f9a8-48b6-8815-43aae5131b94",
...

```

In this case, the id for `esxi-2.vrack.vsphere.local` is `b318fe37-f9a8-48b6-8815-43aae5131b94`.

5. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
6. At the end of the file, add the following line:

```
esx.upgrade.skip.host.ids=hostid1,hostid2
```

Replace the host ids with the information from step 4. If you are including multiple host ids, do not add any spaces between them. For example:

```
esx.upgrade.skip.host.ids=60927f26-8910-4dd3-8435-8bb7aef5f659,6c516864-
b6de-4537-90e4-c0d711e5befb,65c206aa-2561-420e-8c5c-e51b9843f93d
```

7. Save and close the file.
8. Ensure that the ownership of the `application-prod.properties` file is `vcf_lcm:vcf`.
9. Restart the LCM server by typing the following command in the console window:

```
systemctl restart lcm
```

The hosts added to the `application-prod.properties` are not updated when you update the workload domain.

Upgrade ESXi with Custom ISOs

For clusters in workload domains with vSphere Lifecycle Manager baselines, you can upgrade ESXi with a custom ISO from your vendor. VMware Cloud Foundation 4.4.1.1 and later support multiple custom ISOs in a single ESXi upgrade in cases where specific clusters or workload domains require different custom ISOs.

Download the appropriate vendor-specific ISOs on a computer with internet access. If no vendor-specific ISO is available for the required version of ESXi, then you can create one. See [Create a Custom ISO Image for ESXi](#).

1. Download the VMware Software Update bundle for VMware ESXi. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).

To use an async patch version of ESXi, enable the patch with the Async Patch Tool before proceeding to the next step. See the [Async Patch Tool](#).

2. Using SSH, log in to the SDDC Manager appliance.
3. Create a directory for the vendor ISO(s) under the `/nfs/vmware/vcf/nfs-mount` directory. For example, `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries`.
4. Copy the vendor-specific ISO(s) to the directory you created on the SDDC Manager appliance. For example, you can copy the ISO to the `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries` directory.
5. Change permissions on the directory where you copied the ISO(s). For example,

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/
```

6. Change owner to `vcf`.


```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/
```
7. Create an ESX custom image JSON using the following template.

```
{
  "esxCustomImageSpecList": [{
    "bundleId": "bundle ID of the ESXi bundle you downloaded",
    "targetEsxVersion": "ESXi version for the target VMware Cloud Foundation version",
    "useVcfBundle": false,
    "domainId": "xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx"
```

```

,
"clusterId": "xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx"
,
"customIsoAbsolutePath": "Path_to_custom_ISO"
]]
}

```

where

Parameter	Description and Example Value
bundleId	ID of the ESXi upgrade bundle you downloaded. You can retrieve the bundle ID by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the bundle ID. For example, 8c0de63d-b522-4db8-be6c-f1e0ab7ef554. The bundle ID for an async patch looks slightly different. For example: 5dc57fe6-2c23-49fc-967c-0bea1bfea0f1-apTool. NOTE If an incorrect bundle ID is provided, the upgrade will proceed with the VMware Cloud Foundation stock ISO and replace the custom VIBs in your environment with the stock VIBs.
targetEsxVersion	Version of the ESXi bundle you downloaded. You can retrieve the target ESXi version by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the "Update to Version".
useVcfBundle	Specifies whether the VMware Cloud Foundation ESXi bundle is to be used for the upgrade. NOTE If you want to upgrade with a custom ISO image, ensure that this is set to false .
domainId (optional, VCF 4.4.1.1 and later only)	ID of the specific workload domain for the custom ISO. Use the VMware Cloud Foundation API (GET /v1/domains) to get the IDs for your workload domains.
clusterId (optional, VCF 4.4.1.1 and later only)	ID of the specific cluster within a workload domain to apply the custom ISO. If you do not specify a clusterId , the custom ISO will be applied to all clusters in the workload domain. Use the VMware Cloud Foundation API (GET /v1/clusters) to get the IDs for your clusters.
customIsoAbsolutePath	Path to the custom ISO file on the SDDC Manager appliance. For example, /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/VMware-VMvisor-Installer-7.0.0.update01-17325551.x86_64-DellEMC_Customized-A01.iso

Here is an example of a completed JSON template.

```

{
"esxCustomImageSpecList": [{
"bundleId": "8c0de63d-b522-4db8-be6c-f1e0ab7ef554",
"targetEsxVersion": "8.0.1-xxxxxxxx",

```

```

"useVcfBundle": false,
"customIsoAbsolutePath":
"/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/VMware-VMvisor-
Installer-8.0.0.update01-xxxxxxxx.x86_64-DellEMC_Customized-A01.iso"
}]
}

```

Here is an example of a completed JSON template with multiple ISOs using a single workload domain and specified clusters (VCF 4.4.1.1 and later only).

```

{
  "esxCustomImageSpecList": [
    {
      "bundleId": "aa7b16b1-d719-44b7-9ced-51bb02ca84f4",
      "targetEsxVersion": "8.0.2-xxxxxxxx",
      "useVcfBundle": false,
      "domainId": "1b7b16b1-d719-44b7-9ced-51bb02ca84b2",
      "clusterId": "c37b16b1-d719-44b7-9ced-51bb02ca84f4",
      "customIsoAbsolutePath": "/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-
binaries/VMware-ESXi-7.0.2-17867351-DELL.zip"
    },
    {
      "bundleId": "aa7b16b1-d719-44b7-9ced-51bb02ca84f4",
      "targetEsxVersion": "7.0.1-18150133",
      "useVcfBundle": false,
      "domainId": "1b7b16b1-d719-44b7-9ced-51bb02ca84b2",
      "customIsoAbsolutePath": "/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-
binaries/VMware-ESXi-7.0.2-17867351-HP.zip"
    }
  ]
}

```

8. Save the JSON file as `esx-custom-image-upgrade-spec.json` in the `/nfs/vmware/vcf/nfs-mount`.

NOTE

If the JSON file is not saved in the correct directory, the stock VMware Cloud Foundation ISO is used for the upgrade and the custom VIBs are overwritten.

9. Set the correct permissions on the `/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json` file:

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

10. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
11. In the `lcm.esx.upgrade.custom.image.spec=` parameter, add the path to the JSON file.
For example, `lcm.esx.upgrade.custom.image.spec=/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json`
12. In the navigation pane, click **Inventory** › **Workload Domains**.
13. On the Workload Domains page, click the domain you are upgrading and then click the **Updates/Patches** tab.
14. Schedule the ESXi upgrade bundle.
15. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).
16. After the upgrade is complete, confirm the ESXi version by clicking **Current Versions**. The ESXi hosts table displays the current ESXi version.

Upgrade ESXi with VMware Cloud Foundation Stock ISO and Async Drivers

For clusters in workload domains with vLCM baselines, you can apply the stock ESXi upgrade bundle with specified async drivers.

Download the appropriate async drivers for your hardware on a computer with internet access.

1. Download the VMware Cloud Foundation ESXi upgrade bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
2. Using SSH, log in to the SDDC Manager appliance.
3. Create a directory for the vendor provided async drivers under the `/nfs/vmware/vcf/nfs-mount` directory. For example, `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers`.
4. Copy the async drivers to the directory you created on the SDDC Manager appliance. For example, you can copy the drivers to the `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers` directory.
5. Change permissions on the directory where you copied the drivers. For example,


```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers
```
6. Change owner to vcf.


```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers
```
7. Create an ESX custom image JSON using the following template.

```
{
  "esxCustomImageSpecList": [{
    "bundleId": "bundle ID of the ESXi bundle you downloaded",
    "targetEsxVersion": "ESXi version for the target VMware Cloud Foundation version",
    "useVcfBundle": true,
    "esxPatchesAbsolutePaths": [Path_to_Drivers]
  }]
}
```

where

Parameter	Description and Example Value
bundleId	ID of the ESXi upgrade bundle you downloaded. You can retrieve the bundle ID by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the bundle ID. For example, <code>8c0de63d-b522-4db8-be6c-f1e0ab7ef554</code> .
targetEsxVersion	Version of the ESXi upgrade bundle you downloaded. You can retrieve the ESXi target version by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the "Update to Version".
useVcfBundle	Specifies whether the ESXi bundle is to be used for the upgrade. Set this to true .
esxPatchesAbsolutePaths	Path to the async drivers on the SDDC Manager appliance. For example, <code>/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers/VMW-ESX-6.7.0-smartpqi-1.0.2.1038-offline_bundle-8984687.zip</code>

Here is an example of a completed JSON template.

```
{
  "esxCustomImageSpecList": [{
    "bundleId": "411bea6a-b26c-4a15-9443-03f453c68752-apTool",
    "targetEsxVersion": "7.0.3-21053776",
    "useVcfBundle": true,
    "esxPatchesAbsolutePaths": ["/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/
drivers/HPE-703.0.0.10.9.5.14-Aug2022-Synergy-Addon-depot.zip"]
  }]
}
```

8. Save the JSON file as `esx-custom-image-upgrade-spec.json` in the `/nfs/vmware/vcf/nfs-mount`.

NOTE

If the JSON file is not saved in the correct directory, the stock VMware Cloud Foundation ISO is used for the upgrade and the custom VIBs are overwritten.

9. Set the correct permissions on the `/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json` file:

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

10. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.

11. In the `lcm.esx.upgrade.custom.image.spec=` parameter, add the path to the JSON file.

For example, `lcm.esx.upgrade.custom.image.spec=/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json`

12. In the navigation pane, click **Inventory > Workload Domains**.
13. On the Workload Domain page, click the management domain.

14. On the Domain Summary page, click the **Updates/Patches** tab.
15. In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update bundle for VMware ESXi.
16. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).
17. After the upgrade is complete, confirm the ESXi version by clicking **Current Versions**. The ESXi hosts table displays the current ESXi version.

Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2

Prior to VMware Cloud Foundation 5.2.1, workload domains can use either vSphere Lifecycle Manager baselines or vSphere Lifecycle Manager images for ESXi host upgrade. The following procedure describes upgrading ESXi hosts in workload domains that use vSphere Lifecycle Manager images when your target version is VMware Cloud Foundation 5.2.

- Validate that the ESXi passwords are valid.
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.
- You must upgrade NSX and vCenter Server before you can upgrade ESXi hosts with a vSphere Lifecycle Manager image.
- If you want to add firmware to the vSphere Lifecycle Manager image, you must install the Hardware Support Manager from your vendor. See [Firmware Updates](#).

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager baselines when your target version is VMware Cloud Foundation 5.2, see [Upgrade ESXi with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation](#).

VMware Cloud Foundation 5.2.1 and later support workload domains that include vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters. If you are upgrading to VMware Cloud Foundation 5.2.1, see [Upgrade ESXi for VMware Cloud Foundation 5.2.1](#).

You create a vSphere Lifecycle Manager image for upgrading ESXi hosts using the vSphere Client. During the creation of the image, you define the ESXi version and can optionally add vendor add-ons, components, and firmware. After you extract the vSphere Lifecycle Manager image into SDDC Manager, the ESXi update will be available for the relevant VI workload domains.

1. Log in to the management domain vCenter Server using the vSphere Client.
2. Create a vSphere Lifecycle Manager image.
 - a) Right-click the management domain data center and select **New Cluster**.
 - b) Enter a name for the cluster (for example, `ESXi upgrade image`) and click **Next**.
Keep the default settings for everything except the cluster name.

- c) Define the vSphere Lifecycle manager image and click **Next**.

Image Element	Description
ESXi Version	From the ESXi Version drop-down menu, select the ESXi version specified in the VMware Cloud Foundation BOM. If the ESXi version does not appear in the drop-down menu, see Working With the vSphere Lifecycle Manager Depot .
Vendor Add-On (optional)	To add a vendor add-on to the image, click Select and select a vendor add-on.

You can customize the image components, firmware, and drivers later.

- d) Click **Finish**.
- e) After the cluster is created successfully, click the **Updates** tab for the new cluster to further customize it, if needed.
- f) Click **Hosts > Image** and then click **Edit**.

- g) Edit the vSphere Lifecycle manager image properties and click **Save**.

You already specified the ESXi version and optional vendor add-on, but you can modify those settings as required.

Image Element	Description
ESXi Version	From the ESXi Version drop-down menu, select the ESXi version specified in the VMware Cloud Foundation BOM. If the ESXi version does not appear in the drop-down menu, see Synchronize the vSphere Lifecycle Manager Depot and Import Updates to the vSphere Lifecycle Manager Depot .
Vendor Add-On (optional)	To add a vendor add-on to the image, click Select and select a vendor add-on.
Firmware and Drivers Add-On (optional)	To add a firmware add-on to the image, click Select . In the Select Firmware and Drivers Addon dialog box, specify a hardware support manager and select a firmware add-on to add to the image. Selecting a firmware add-on for a family of vendor servers is possible only if the respective vendor-provided hardware support manager is registered as an extension to the vCenter Server where vSphere Lifecycle Manager runs.
Components	To add components to the image: <ul style="list-style-type: none"> • Click Show details. • Click Add Components. • Select the components and their corresponding versions to add to the image.

vSphere saves the cluster image.

3. Extract the vSphere Lifecycle Manager image into SDDC Manager.
 - a) In the SDDC Manager UI, click **Lifecycle Management > Image Management** .
 - b) Click **Import Image**.
 - c) In the Option 1 section, select the management domain from the drop-down menu.
 - d) In the Cluster drop-down, select the cluster from which you want to extract the vSphere Lifecycle manager image. For example, *ESXi upgrade image*.

Image Management

Available Images Import Image

Make a cluster image available to Cloud Foundation by either extracting or importing an image.


Option 1 Extract a Cluster Image

Extract a cluster image assigned to a cluster that was updated in vCenter.

Select Workload Domain

Select Cluster

Cluster Image Name

 EXTRACT CLUSTER IMAGE

e) Enter a name for the cluster image and click **Extract Cluster Image**.

You can view status in the **Tasks** panel.

4. Upgrade ESXi hosts with the vSphere Lifecycle Manager image.





- Navigate to the **Updates** tab of the VI workload domain.
- In the Available Updates section, click **Configure Update**.
- Click **Next**.
- Select the clusters to upgrade and click **Next**.

The default setting is to upgrade all clusters. To upgrade specific clusters, click **Enable cluster-level selection** and select the clusters to upgrade.

- Select the cluster and the cluster image, and click **Apply Image**.
- Click **Next**.
- Select the upgrade options and click **Next**.

Upgrade Options

Select options to optimize your upgrade experience.

- Enable sequential cluster upgrade 
- Enable Quick Boot 
- Enforce Live Patch 
- Migrate Powered Off and Suspended VMs 

By default, the selected clusters are upgraded in parallel. If you selected more than five clusters to be upgraded, the first five are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Select **Enable Quick Boot** to reduce the upgrade time by skipping the physical reboot of the host.

Select **Enforce Live Patch** when the cluster image includes a Live Patch. With the Enforce Live Patch option, vSphere Lifecycle Manager does not place the hosts in the cluster into maintenance mode, hosts are not rebooted, and there is no need to migrate the virtual machines running on the hosts in the cluster.

Select **Migrate Powered Off and Suspended VMs** to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster.

h) Review the settings, and click **Finish**.

VMware Cloud Foundation runs a cluster image hardware compatibility and compliance check. Resolve any reported issues before proceeding.

i) Click **Schedule Update** and click **Next**.

j) Select **Upgrade Now** or **Schedule Update** and click **Finish**.

k) Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [Upgrade vSAN on-disk format versions](#).

Firmware Updates

You can use vSphere Lifecycle Manager images to perform firmware updates on the ESXi hosts in a cluster. Using a vSphere Lifecycle Manager image simplifies the host update operation. With a single operation, you update both the software and the firmware on the host.

To apply firmware updates to hosts in a cluster, you must deploy and configure a vendor provided software module called hardware support manager. The deployment method and the management of a hardware support manager is determined by the respective OEM. For example, the hardware support manager that Dell EMC provides is part of their host management solution, OpenManage Integration for VMware vCenter (OMIVV), which you deploy as an appliance. See [Deploying Hardware Support Managers](#).

You must deploy the hardware support manager appliance on a host with sufficient disk space. After you deploy the appliance, you must power on the appliance virtual machine, log in to the appliance as an administrator, and register the appliance as a vCenter Server extension. Each hardware support manager has its own mechanism of managing firmware packages and making firmware add-ons available for you to choose.

For detailed information about deploying, configuring, and managing hardware support managers, refer to the vendor-provided documentation.

Update License Keys for a Workload Domain

If upgrading from a VMware Cloud Foundation version prior to 5.0, you need to update your license keys to support vSAN 8.x and vSphere 8.x.

You need a new license key for vSAN 8.x and vSphere 8.x. Prior to VMware Cloud Foundation 5.1.1, you must add and update the component license key for each upgraded component in the SDDC Manager UI as described below.

With VMware Cloud Foundation 5.1.1 and later, you can add a component license key as described below, or add a solution license key in the vSphere Client. See [Managing vSphere Licenses](#) for information about using a solution license key for VMware ESXi and vCenter Server. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

You first add the new component license key to SDDC Manager. This must be done once per license instance. You then apply the license key to the component on a per workload domain basis.

1. Add a new component license key to the SDDC Manager inventory.
 - a) In the navigation pane, click **Administration** › **Licensing**.
 - b) On the **Licensing** page, click **+ License Key**.
 - c) Select a product from the drop-down menu.
 - d) Enter the license key.
 - e) Enter a description for the license key.
 - f) Click **Add**.
 - g) Repeat for each license key to be added.
2. Update a license key for a workload domain component.
 - a) In the navigation pane, click **Inventory** › **Workload Domains**.
 - b) On the **Workload Domains** page, click the domain you are upgrading.
 - c) On the **Summary** tab, expand the red error banner, and click **Update Licenses**.
 - d) On the **Update Licenses** page, click **Next**.
 - e) Select the products to update and click **Next**.
 - f) For each product, select a new license key from the list, and select the entity to which the licensekey should be applied and click **Next**.
 - g) On the Review pane, review each license key and click **Submit**.
The new license keys will be applied to the workload domain. Monitor the task in the **Tasks** pane in SDDC Manager.

Upgrade vSphere Distributed Switch versions

[Optional] Upgrade the distributed switch to take advantage of features that are available only in the later versions.

ESXi and vCenter Upgrades are completed.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Upgrade** › **Upgrade Distributed Switch**
3. Select the vSphere Distributed Switch version that you want to upgrade the switch to and click **Next**

The vSphere Distributed Switch is successfully upgraded.

Upgrade vSAN on-disk format versions

[Optional] Upgrade the vSAN on-disk format version to take advantage of features that are available only in the later versions.

- ESXi and vCenter Upgrades are completed
- Verify that the disks are in a healthy state. Navigate to the Disk Management page to verify the object status.
- Verify that your hosts are not in maintenance mode. When upgrading the disk format, do not place the hosts in maintenance mode.
- Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster. For information about vSAN resynchronization, see vSphere Monitoring and Performance
- The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure.

1. Navigate to the vSAN cluster.

2. Click the **Configure** tab.
3. Under **vSAN**, select **Disk Management**.
4. Click **Pre-check Upgrade**. The upgrade pre-check analyzes the cluster to uncover any issues that might prevent a successful upgrade. Some of the items checked are host status, disk status, network status, and object status. Upgrade issues are displayed in the **Disk pre-check status text** box.
5. Click **Upgrade**.
6. Click **Yes** on the Upgrade dialog box to perform the upgrade of the on-disk format.

vSAN successfully upgrades the on-disk format. The On-disk Format Version column displays the disk format version of storage devices in the cluster

Upgrade VI Workload Domains to VMware Cloud Foundation 5.2.x

The management domain in your environment must be upgraded before you upgrade VI workload domains. To upgrade to VMware Cloud Foundation 5.2.x, all VI workload domains in your environment must be at VMware Cloud Foundation 4.5 or higher. If your environment is at a version lower than 4.5, you must upgrade the workload domains to 4.5 and then upgrade to 5.2.x.

Within a VI workload domain, components must be upgraded in the following order.

1. NSX.
2. vCenter Server.
3. ESXi.
4. Workload Management on clusters that have vSphere with Tanzu. Workload Management can be upgraded through vCenter Server. See [Updating the vSphere with Tanzu Environment](#).
5. If you suppressed the Enter Maintenance Mode prechecks for ESXi or NSX, delete the following lines from the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file and restart the LCM service:
`lcm.nsxt.suppress.dry.run.emm.check=true`

`lcm.esx.suppress.dry.run.emm.check.failures=true`
6. If you have stretched clusters in your environment, upgrade the vSAN witness host. See [Upgrade vSAN Witness Host for VMware Cloud Foundation](#).
7. For NFS-based workload domains, add a static route for hosts to access NFS storage over the NFS gateway. See [Post Upgrade Steps for NFS-Based VI Workload Domains](#).

After all upgrades have completed successfully:

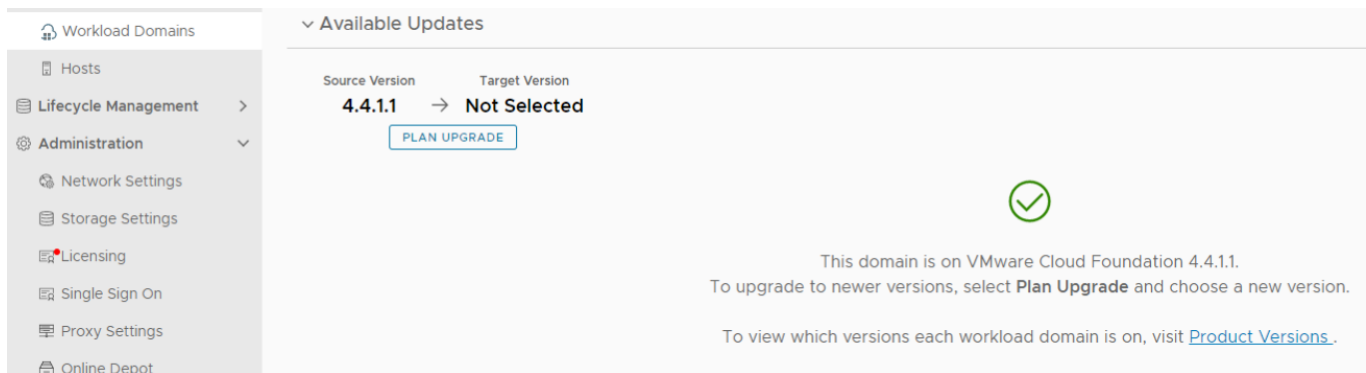
1. Remove the VM snapshots you took before starting the update.
2. Take a backup of the newly installed components.

Plan VI Workload Domain Upgrade

Before proceeding with a VI workload domain upgrade you must first plan the upgrade to your target version.

[Upgrade the Management Domain to VMware Cloud Foundation 5.2.x](#).

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. On the **Workload Domains** page, click the workload domain you want to upgrade and click the **Updates** tab.
3. Under **Available Updates**, click **PLAN UPGRADE**.



Workload Domains

- Hosts
- Lifecycle Management >
- Administration ▾
 - Network Settings
 - Storage Settings
 - Licensing
 - Single Sign On
 - Proxy Settings
 - Online Depot

Available Updates

Source Version: 4.4.1.1 → Target Version: Not Selected

PLAN UPGRADE

✓

This domain is on VMware Cloud Foundation 4.4.1.1.
To upgrade to newer versions, select **Plan Upgrade** and choose a new version.
To view which versions each workload domain is on, visit [Product Versions](#).

4. On the **Plan Upgrade for VMware Cloud Foundation** screen, select the target version from the drop-down, and click **CONFIRM**.

CAUTION

You must upgrade all VI workload domains to VMware Cloud Foundation 5.x. Upgrading to a higher 4.x release once the management domain has been upgraded to 5.x is unsupported.

Plan Upgrade for VMware Cloud Foundation ?

Select a version of VMware Cloud Foundation to upgrade this workload domain. Only versions up to the management workload domain's version are available.

i To activate or deactivate compatibility checks on the target version, see KB 90074.

VMware Cloud Foundation 5.0.0.0 ▼

Select Version

Upgrade Summary i

Software Component	Source Version	Source Build	Target Version	Target Build
vRealize Suite Lifecycle Manager	8.6.2	19221620	8.10.0	21331275
VMware NSX	3.1.3.7.4	19762317	4.1.0.1.0	21500054
VMware ESXi	7.0.3	19482537	8.0.1	21495797
SDDC Manager	4.4.1.1	19948546	5.0.0.0	21595347
VMware vCenter Server Appliance	7.0.3.00500	19480866	8.0.1.00000	21560480

By confirming, you are selecting a target version for this workload domain which will override your current target version and upgrade path. No upgrades will be performed until you schedule an upgrade. You can change your target version at any time until an upgrade is run or scheduled. For more information, review

CANCEL

CONFIRM

Bundles applicable to the chosen release will be made available to the VI workload domain.

Available Updates

Source Version Target Version
4.4.1.1 → **5.0.0.0**

PLAN UPGRADE

○ Refreshing bundles for Cloud Foundation 5.0.0.0...

Perform Update Precheck in SDDC Manager

You must perform a precheck in SDDC Manager before applying an update bundle to ensure that your environment is ready for the update.

Bundle-level pre-checks for vCenter are available in VMware Cloud Foundation.

NOTE

Because ESXi bundle-level pre-checks only work in minor-version upgrades (for example: from ESXi 7.x through 7.y, or from ESXi 8.x through 8.y), these prechecks do not run in VMware Cloud Foundation.

If you silence a vSAN Skyline Health alert in the vSphere Client, SDDC Manager skips the related precheck and indicates which precheck it skipped. Click **RESTORE PRECHECK** to include the silenced precheck. For example: You can also silence failed vSAN prechecks in the SDDC Manager UI by clicking **Silence Precheck**. Silenced prechecks do not trigger warnings or block upgrades.

IMPORTANT

Only silence alerts if you know that they are incorrect. Do not silence alerts for real issues that require remediation.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the workload domain where you want to run the precheck.
3. On the domain summary page, click the **Updates** tab.

(The following image is a sample screenshot and may not reflect current product versions.)



NOTE

It is recommended that you Precheck your workload domain prior to performing an upgrade.

4. Click **RUN PRECHECK** to select the components in the workload domain you want to precheck.
 - a) You can select to run a Precheck only on vCenter or the vSphere cluster. All components in the workload domain are selected by default. To perform a precheck on certain components, choose **Custom selection**.
 - b) If there are pending upgrade bundles available, then the "Target Version" dropdown contains "General Upgrade Readiness" and the available VMware Cloud Foundation versions to upgrade to. If there is an available VMware Cloud Foundation upgrade version, there will be extra checks - bundle-level prechecks for hosts, vCenter Server, and so forth. The version specific prechecks will only run prechecks on components that have available upgrade bundles downloaded.
5. When the precheck begins, a progress message appears indicating the precheck progress and the time when the precheck began.

NOTE

Parallel precheck workflows are supported. If you want to precheck multiple domains, you can repeat steps 1-5 for each of them without waiting for step 5 to finish.

6. Once the Precheck is complete, the report appears. Click through **ALL**, **ERRORS**, **WARNINGS**, and **SILENCED** to filter and browse through the results.

7. To see details for a task, click the expander arrow.

If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **RETRY ALL FAILED RESOURCES** to retry all failed tasks.

8. If ESXi hosts display a driver incompatibility issue when updating a VI workload domain using vSphere Lifecycle Manager baselines, perform the following steps:
 1. Identify the controller with the HCL issue.
 2. For the given controller, identify the supported driver and firmware versions on the source and target ESXi versions.
 3. Upgrade the firmware, if required.
 4. Upgrade the driver manually on the ESXi host and retry the task at which the upgrade failed.

9. If the workload domain contains a host that includes pinned VMs, the precheck fails at the Enter Maintenance Mode step. If the host can enter maintenance mode through vCenter Server UI, you can suppress this check for NSX and ESXi in VMware Cloud Foundation by following the steps below.
 1. Log in to SDDC Manager by using a Secure Shell (SSH) client with the user name `vcf` and password.
 2. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
 3. Add the following line to the end of the file:


```
lcm.nsxt.suppress.dry.run.emm.check=true
```

```
lcm.esx.suppress.dry.run.emm.check.failures=true
```
 4. Restart Lifecycle Management by typing the following command in the console window.


```
systemctl restart lcm
```
 5. After Lifecycle Management is restarted, run the precheck again.

The precheck result is displayed at the top of the Upgrade Precheck Details window. If you click **Exit Details**, the precheck result is displayed at the top of the Precheck section in the Updates tab.

Ensure that the precheck results are green before proceeding. Although a failed precheck will not prevent the upgrade from proceeding, it may cause the update to fail.

Upgrade NSX for VMware Cloud Foundation in a Federated Environment

If NSX Federation is configured between two VMware Cloud Foundation instances, SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must manually upgrade the NSX Global Managers for each instance.

Download NSX Global Manager Upgrade Bundle

SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must download the NSX upgrade bundle manually to upgrade the NSX Global Managers.

1. Log in to the Broadcom Support Portal and browse to **My Downloads > VMware NSX**.
2. Click the version of NSX to which you are upgrading.
3. Locate the **NSX version Upgrade Bundle** and verify that the upgrade bundle filename extension ends with `.mub`.

The upgrade bundle filename has the following format `VMware-NSX-upgrade-bundle-versionnumber.buildnumber.mub`.

4. Click the download icon to download the upgrade bundle to the system where you access the NSX Global Manager UI.

Upgrade the Upgrade Coordinator for NSX Federation

The upgrade coordinator runs in the NSX Manager. It is a self-contained web application that orchestrates the upgrade process of hosts, NSX Edge cluster, NSX Controller cluster, and the management plane.

The upgrade coordinator guides you through the upgrade sequence. You can track the upgrade process and, if necessary, you can pause and resume the upgrade process from the UI.

1. In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Proceed to Upgrade**.
4. Navigate to the upgrade bundle .mub file you downloaded or paste the download URL link.
 - Click **Browse** to navigate to the location you downloaded the upgrade bundle file.
 - Paste the VMware download portal URL where the upgrade bundle .mub file is located.
5. Click **Upload**.
When the file is uploaded, the **Begin Upgrade** button appears.
6. Click **Begin Upgrade** to upgrade the upgrade coordinator.

NOTE

Upgrade one upgrade coordinator at a time.

7. Read and accept the EULA terms and accept the notification to upgrade the upgrade coordinator..
8. Click **Run Pre-Checks** to verify that all NSX components are ready for upgrade.
The pre-check checks for component connectivity, version compatibility, and component status.
9. Resolve any warning notifications to avoid problems during the upgrade.

Upgrade NSX Global Managers for VMware Cloud Foundation

Manually upgrade the NSX Global Managers when NSX Federation is configured between two VMware Cloud Foundation instances.

Before you can upgrade NSX Global Managers, you must upgrade all VMware Cloud Foundation instances in the NSX Federation, including NSX Local Managers, using SDDC Manager.

1. In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Start** to upgrade the management plane and then click **Accept**.
4. On the Select Upgrade Plan page, select **Plan Your Upgrade** and click **Next**.

The NSX Manager UI, API, and CLI are not accessible until the upgrade finishes and the management plane is restarted.

Upgrade NSX for VMware Cloud Foundation 5.2.x

Upgrade NSX in the management domain and VI workload domains. VMware Cloud Foundation 5.2.1 supports in-place host upgrades for clusters that use vSphere Lifecycle Manager baselines.

Until SDDC Manager is upgraded to version 5.2, you must upgrade NSX in the management domain before you upgrade NSX in a VI workload domain. Once SDDC Manager is at version 5.2 or later, you can upgrade NSX in VI workload domains before or after upgrading NSX in the management domain.

Upgrading NSX involves the following components:

- Upgrade Coordinator
- NSX Edges/Clusters (if deployed)
- Host clusters
- NSX Manager cluster

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates/Patches** tab.

When you upgrade NSX components for a selected VI workload domain, those components are upgraded for all VI workload domains that share the NSX Manager cluster.

3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

NOTE

The NSX precheck runs on all VI workload domains in your environment that share the NSX Manager cluster.

4. For VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for NSX.
 - b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.
 By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.
 - c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.
 By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.
 By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) If you selected the **Schedule Upgrade** option, specify the date and time for the NSX bundle to be applied and click **Next**.
- f) On the Review page, review your settings and click **Finish**.

If you selected **Upgrade Now**, the NSX upgrade begins and the upgrade components are displayed. The upgrade view displayed here pertains to the workload domain where you applied the bundle. Click the link to the associated workload domains to see the components pertaining to those workload domains. If you selected **Schedule Upgrade**, the upgrade begins at the time and date you specified.

5. For VMware Cloud Foundation 5.2.1:

- a) In the Available Updates section, click the **Configure Update** button.
- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.
By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.
- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.
By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.
By default ESXi hosts are placed into maintenance mode during an upgrade. Starting with VMware Cloud Foundation 5.2.1, in-place upgrades are available for workload domains in which all the clusters use vSphere Lifecycle Manager baselines. If NSX Manager is shared between workload domains, in-place upgrade is only available if all the clusters in all the workload domains that share the NSX Manager use vLCM baselines. If the option is available, you can select **In-place** as the upgrade mode to avoid powering off and placing hosts into maintenance mode before the upgrade.

NOTE

To perform an in-place upgrade, the target NSX version must be the VMware Cloud Foundation 5.2.1 BOM version or later.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) On the Review page, review your settings and click **Run Precheck**.
The precheck begins. Resolve any issues until the precheck succeeds.
- f) After the precheck succeeds, click **Schedule Update** and select an option.

6. Monitor the upgrade progress. See [monitor-update.dita](#).

If a component upgrade fails, the failure is displayed across all associated workload domains. Resolve the issue and retry the failed task.

When all NSX workload components are upgraded successfully, a message with a green background and check mark is displayed.

Upgrade vCenter Server for VMware Cloud Foundation 5.2.x

The upgrade bundle for VMware vCenter Server is used to upgrade the vCenter Server instances managed by SDDC Manager. Upgrade vCenter Server in the management domain before upgrading vCenter Server in VI workload domains.

- Download the VMware vCenter Server upgrade bundle. See [downloading-vmware-cloud-foundation-bundles.dita](#).
- Take a file-based backup of the vCenter Server appliance before starting the upgrade. See [manually-back-up-vcenter-server.dita](#).

NOTE

After taking a backup, do not make any changes to the vCenter Server inventory or settings until the upgrade completes successfully.

- If your workload domain contains Workload Management (vSphere with Tanzu) enabled clusters, the supported target release depends on the version of Kubernetes (K8s) currently running in the cluster. Older versions of K8s might require a specific upgrade sequence. See [KB 92227](#) for more information.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

4. Upgrading to VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for vCenter Server.
 - b) Click **Confirm** to confirm that you have taken a file-based backup of the vCenter Server appliance before starting the upgrade.
 - c) If you selected **Schedule Update**, click the date and time for the bundle to be applied and click **Schedule**.
 - d) If you are upgrading from VMware Cloud Foundation 4.5.x, enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.
 - e) Review the upgrade settings and click **Finish**.
5. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 5.x:
 - a) In the Available Updates section, click **Configure Update**.
 - b) Select the upgrade mechanism and click **Next**.

Option	Description
vCenter Reduced Downtime Upgrade	<p>The reduced downtime upgrade process uses a migration-based approach. In this approach, a new vCenter Server Appliance is deployed and the current vCenter data and configuration is copied to it.</p> <p>During the preparation phase of a reduced downtime upgrade, the source vCenter Server Appliance and all resources remain online. The only downtime occurs when the source vCenter Server Appliance is stopped, the configuration is switched over to the target vCenter, and the services are started. The downtime is expected to take</p>

Table continued on next page

Continued from previous page

Option	Description
	approximately 5 minutes under ideal network, CPU, memory, and storage provisioning. NOTE To perform a vCenter Reduced Downtime Upgrade, the target vCenter version must be the VMware Cloud Foundation 5.2.1 BOM version or later.
vCenter Regular Upgrade	During a regular upgrade, the vCenter Server Appliance is offline for the duration of the upgrade.

- c) Select a backup option and click **Next**.
 d) For an RDU update, provide a temporary network to be used only during the upgrade and click **Next**.

Automatic	Automatically assign network information.
Static	Enter an IP address, subnet mask, and gateway. The IP address must be in the management subnet.

- e) Schedule the update and click **Next**.

For vCenter Reduced Downtime Upgrade	Select scheduling options for the preparation and switchover phases of the upgrade. NOTE If you are scheduling the switchover phase, you must allow a minimum of 4 hours between the start of preparation and the start of switchover.
For vCenter Regular Upgrade	Select an Upgrade Now or Schedule Update .

- f) Review the upgrade settings and click **Finish**.
6. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 4.5.x:
- In the Available Updates section, click **Configure Update**.
 - Enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.
 - Select a backup option and click **Next**.
 - Schedule the update and click **Next**.
 - Review the upgrade settings and click **Finish**.
7. Monitor the upgrade progress. See [monitor-update.dita](#).
8. After the upgrade is complete, remove the old vCenter Server appliance (if applicable).

NOTE

Removing the old vCenter is only required for major upgrades. If you performed a vCenter RDU patch upgrade, the old vCenter is automatically removed after a successful upgrade.

If the upgrade fails, resolve the issue and retry the failed task. If you cannot resolve the issue, restore vCenter Server using the file-based backup. See [restore-vcenter-server.dita](#). vCenter RDU upgrades perform automatic rollback if the upgrade fails.

Once the upgrade successfully completes, use the vSphere Client to change the vSphere DRS Automation Level setting back to the original value (before you took a file-based backup) for each vSphere cluster that is managed by the vCenter Server. See [KB 87631](#) for information about using VMware PowerCLI to change the vSphere DRS Automation Level.

Upgrade ESXi for VMware Cloud Foundation 5.2.1

VMware Cloud Foundation 5.2.1 and later support workload domains that include vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters. There is a single procedure for upgrading both vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters.

- Validate that the ESXi passwords are valid.
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.
- For clusters that use vSphere Lifecycle Manager images:
 - You must upgrade NSX and vCenter Server before you can upgrade ESXi hosts with a vSphere Lifecycle Manager image.
 - If you want to add firmware to the vSphere Lifecycle Manager image, you must install the Hardware Support Manager from your vendor. See [../shared-content/topics/adding-firmware-and-components-to-a-cluster-image.dita#GUID-198C2DE5-8DC1-4327-9914-C4677F4964D5-en](#).
 - A supported vSphere Lifecycle Manager image must be available in SDDC Manager. See steps 1-3 in [../shared-content/topics/upgrade-esxi-with-vsphere-lifecycle-manager-images-for-vmware-cloud-foundation.dita](#) for more information.
- For clusters that use vSphere Lifecycle Manager baselines, download the ESXi bundle. See [../shared-content/topics/downloading-vmware-cloud-foundation-bundles.dita#GUID-4D553D24-9DBA-47C6-A4FE-D737329C1C26-en](#).

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager images when your target version is VMware Cloud Foundation 5.2, see [../shared-content/topics/upgrade-esxi-with-vsphere-lifecycle-manager-images-for-vmware-cloud-foundation.dita#GUID-A9C387B5-257A-48F4-8923-A64AC65B610F-en](#).

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager baselines when your target version is VMware Cloud Foundation 5.2, see [upgrade-esxi-with-vsphere-lifecycle-manager-baselines-for-vmware-cloud-foundation-5-2.dita](#).

If you are using external (non-vSAN) storage, the following procedure updates the ESXi hosts attached to the external storage. However, updating and patching the storage software and drivers is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability after an ESXi upgrade, consult the vSphere HCL and your storage vendor.

For clusters that use vSphere Lifecycle Manager baselines:

- If you want to skip any hosts while applying an ESXi update a workload domain, you must add these hosts to the `application-prod.properties` file before you begin the update. See "Skip Hosts During ESXi Update".
- To perform ESXi upgrades with custom ISO images or async drivers see "Upgrade ESXi with Custom ISOs" and "Upgrade ESXi with Stock ISO and Async Drivers".

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

4. In the Available Updates section, click **Configure Update**.
5. Read the introductory information and click **Next**.
6. If any clusters in the workload domain use vSphere Lifecycle Manager images, select the clusters to update and click **Next**.

7. Assign an image to each cluster that uses vSphere Lifecycle Manager images and click **Next**.
8. If any clusters in the workload domain use vSphere Lifecycle Manager baselines, select the clusters to upgrade and click **Next**.

The default setting is to upgrade all clusters. To upgrade specific clusters, select **Custom selection** and select the clusters to upgrade.

9. If the workload domain you are upgrading only includes clusters that use vSphere Lifecycle Manager baselines, select a scheduling option.
10. Select the upgrade options and click **Next**.
By default, the selected clusters are upgraded in parallel. If you selected more than ten clusters to be upgraded, the first ten are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Select **Enable Quick Boot** to reduce the upgrade time by skipping the physical reboot of the host.

Select **Migrate Powered Off and Suspended VMs** to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster.

For clusters that use vSphere Lifecycle Manager images, select **Enforce Live Patch** when the cluster image includes a Live Patch. With the **Enforce Live Patch** option, vSphere Lifecycle Manager does not place the hosts in the cluster into maintenance mode, hosts are not rebooted, and there is no need to migrate the virtual machines running on the hosts in the cluster.

11. Review the settings, and click **Finish** or **Run Precheck**.
If the upgrade includes any clusters that use vSphere Lifecycle Manager images VMware Cloud Foundation runs a cluster image hardware compatibility and compliance precheck. Resolve any reported issues before proceeding.
12. After the precheck succeeds, click **Schedule Update**, select a scheduling option, and click **Finish**.
13. Monitor the upgrade progress. See [../shared-content/topics/monitor-update.dita#GUID-68328710-37D1-48DE-B5AB-7120322DB68A-en](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [../shared-content/topics/upgrade-vsan-on-disk-format-versions.dita](#).

Upgrade ESXi with vSphere Lifecycle Manager Baselines for VMware Cloud Foundation

Workload domains can use vSphere Lifecycle Manager baselines or vSphere Lifecycle Manager images. The following procedure describes upgrading ESXi hosts in workload domains that use vSphere Lifecycle Manager baselines.

- Validate that the ESXi passwords are valid.
- Download the ESXi bundle. See [downloading-vmware-cloud-foundation-bundles.dita](#).
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.

For information about upgrading ESXi in VI workload domains that use vSphere Lifecycle Manager images, see [upgrade-esxi-with-vsphere-lifecycle-manager-images-for-vmware-cloud-foundation.dita](#).

By default, the upgrade process upgrades the ESXi hosts in all clusters in a workload domain in parallel. If you have multiple clusters in a workload domain, you can select the clusters to upgrade.

If you want to skip any hosts while applying an ESXi update a workload domain, you must add these hosts to the `application-prod.properties` file before you begin the update. See "Skip Hosts During ESXi Update".

To perform ESXi upgrades with custom ISO images or async drivers see "Upgrade ESXi with Custom ISOs" and "Upgrade ESXi with Stock ISO and Async Drivers".

If you are using external (non-vSAN) storage, the following procedure updates the ESXi hosts attached to the external storage. However, updating and patching the storage software and drivers is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability after an ESXi upgrade, consult the vSphere HCL and your storage vendor.

1. Navigate to the **Updates/Patches** tab of the workload domain.
2. Click **Precheck** to run the upgrade precheck.
Resolve any issues before proceeding with the upgrade.
3. In the Available Updates section, select the target release.
4. Click **Upgrade Now** or **Schedule Update**.
5. If you selected **Schedule Update**, specify the date and time for the bundle to be applied.
6. Select the clusters to upgrade and click **Next**.
The default setting is to upgrade all clusters. To upgrade specific clusters, click **Enable cluster-level selection** and select the clusters to upgrade.
7. Click **Next**.
8. Select the appropriate upgrade options and click **Finish**.
By default, the selected clusters are upgraded in parallel. If you selected more than ten clusters to be upgraded, the first ten are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Click **Enable Quick Boot** if desired. Quick Boot for ESXi hosts is an option that allows vSphere Lifecycle Manager to reduce the upgrade time by skipping the physical reboot of the host.
9. Monitor the upgrade progress. See [monitor-update.dita](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [upgrade-vsan-on-disk-format-versions.dita](#).

Upgrade vSAN Witness Host for VMware Cloud Foundation

If your VMware Cloud Foundation environment contains stretched clusters, update and remediate the vSAN witness host.

Download the ESXi ISO that matches the version listed in the the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes*.

1. In a web browser, log in to vCenter Server at https://vcenter_server_fqdn/ui.
2. Upload the ESXi ISO image file to vSphere Lifecycle Manager.
 - a) Click **Menu** > **Lifecycle Manager**.
 - b) Click the **Imported ISOs** tab.
 - c) Click **Import ISO** and then click **Browse**.
 - d) Navigate to the ESXi ISO file you downloaded and click **Open**.
 - e) After the file is imported, click **Close**.
3. Create a baseline for the ESXi image.
 - a) On the Imported ISOs tab, select the ISO file that you imported, and click **New baseline**.
 - b) Enter a name for the baseline and specify the **Content Type** as Upgrade.
 - c) Click **Next**.
 - d) Select the ISO file you had imported and click **Next**.

- e) Review the details and click **Finish**.
4. Attach the baseline to the vSAN witness host.
 - a) Click **Menu › Hosts and Clusters**.
 - b) In the Inventory panel, click **vCenter › Datacenter**.
 - c) Select the vSAN witness host and click the **Updates** tab.
 - d) Under Attached Baselines, click **Attach › Attach Baseline or Baseline Group**.
 - e) Select the baseline that you had created in step 3 and click **Attach**.
 - f) Click **Check Compliance**.

After the compliance check is completed, the **Status** column for the baseline is displayed as Non-Compliant.

5. Remediate the vSAN witness host and update the ESXi hosts that it contains.
 - a) Right-click the vSAN witness and click **Maintenance Mode › Enter Maintenance Mode**.
 - b) Click **OK**.
 - c) Click the **Updates** tab.
 - d) Select the baseline that you had created in step 3 and click **Remediate**.
 - e) In the End user license agreement dialog box, select the check box and click **OK**.
 - f) In the Remediate dialog box, select the vSAN witness host, and click **Remediate**.

The remediation process might take several minutes. After the remediation is completed, the **Status** column for the baseline is displayed as Compliant.

 - g) Right-click the vSAN witness host and click **Maintenance Mode › Exit Maintenance Mode**.
 - h) Click **OK**.

Skip Hosts During ESXi Update

You can skip hosts while applying an ESXi update to a workload domain. The skipped hosts are not updated.

NOTE

You cannot skip hosts that are part of a VI workload domain that is using vSphere Lifecycle Manager images, since these hosts are updated at the cluster-level and not the host-level.

1. Using SSH, log in to the SDDC Manager appliance with the user name `vcsf` and password you specified in the deployment parameter sheet.
2. Type `su` to switch to the root account.
3. Retrieve the host IDs for the hosts you want to skip.

```
curl 'https://SDDC_MANAGER_IP/v1/hosts' -i -u 'username:password' -X GET -H 'Accept: application/json' |jq json_pp
```

Replace the SDDC Manager FQDN, user name, and password with the information for your environment.

4. Copy the ids for the hosts you want to skip from the output. For example:

```
...
    "fqdn" : "esxi-2.vrack.vsphere.local",
    "esxiVersion" : "6.7.0-16075168",
    "id" : "b318fe37-f9a8-48b6-8815-43aae5131b94",
...

```

In this case, the id for `esxi-2.vrack.vsphere.local` is `b318fe37-f9a8-48b6-8815-43aae5131b94`.

5. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
6. At the end of the file, add the following line:

```
esx.upgrade.skip.host.ids=hostid1,hostid2
```

Replace the host ids with the information from step 4. If you are including multiple host ids, do not add any spaces between them. For example:

```
esx.upgrade.skip.host.ids=60927f26-8910-4dd3-8435-8bb7aef5f659,6c516864-
b6de-4537-90e4-c0d711e5befb,65c206aa-2561-420e-8c5c-e51b9843f93d
```

7. Save and close the file.
8. Ensure that the ownership of the `application-prod.properties` file is `vcf_lcm:vcf`.
9. Restart the LCM server by typing the following command in the console window:

```
systemctl restart lcm
```

The hosts added to the `application-prod.properties` are not updated when you update the workload domain.

Upgrade ESXi with Custom ISOs

For clusters in workload domains with vSphere Lifecycle Manager baselines, you can upgrade ESXi with a custom ISO from your vendor. VMware Cloud Foundation 4.4.1.1 and later support multiple custom ISOs in a single ESXi upgrade in cases where specific clusters or workload domains require different custom ISOs.

Download the appropriate vendor-specific ISOs on a computer with internet access. If no vendor-specific ISO is available for the required version of ESXi, then you can create one. See [create-a-custom-iso-image-for-esxi.dita](#).

1. Download the VMware Software Update bundle for VMware ESXi. See [downloading-vmware-cloud-foundation-bundles.dita](#).

To use an async patch version of ESXi, enable the patch with the Async Patch Tool before proceeding to the next step. See the [ap-tool-5-2.dita](#).

2. Using SSH, log in to the SDDC Manager appliance.
3. Create a directory for the vendor ISO(s) under the `/nfs/vmware/vcf/nfs-mount` directory. For example, `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries`.
4. Copy the vendor-specific ISO(s) to the directory you created on the SDDC Manager appliance. For example, you can copy the ISO to the `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries` directory.
5. Change permissions on the directory where you copied the ISO(s). For example,

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/
```

6. Change owner to `vcf`.

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/
```
7. Create an ESX custom image JSON using the following template.

```
{
  "esxCustomImageSpecList": [{
    "bundleId": "bundle ID of the ESXi bundle you downloaded",
```

```

"targetEsxVersion": "ESXi version for the target VMware Cloud Foundation version",
"useVcfBundle": false,
"domainId": "xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx"
,
"clusterId": "xxxxxxxx-xxxx-xxxx-xxxxxxxxxxxx"
,
"customIsoAbsolutePath": "Path_to_custom_ISO"
}}
}

```

where

Parameter	Description and Example Value
bundleId	ID of the ESXi upgrade bundle you downloaded. You can retrieve the bundle ID by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the bundle ID. For example, 8c0de63d-b522-4db8-be6c-f1e0ab7ef554. The bundle ID for an async patch looks slightly different. For example: 5dc57fe6-2c23-49fc-967c-0bea1bfea0f1-apTool. NOTE If an incorrect bundle ID is provided, the upgrade will proceed with the VMware Cloud Foundation stock ISO and replace the custom VIBs in your environment with the stock VIBs.
targetEsxVersion	Version of the ESXi bundle you downloaded. You can retrieve the target ESXi version by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the "Update to Version".
useVcfBundle	Specifies whether the VMware Cloud Foundation ESXi bundle is to be used for the upgrade. NOTE If you want to upgrade with a custom ISO image, ensure that this is set to false .
domainId (optional, VCF 4.4.1.1 and later only)	ID of the specific workload domain for the custom ISO. Use the VMware Cloud Foundation API (GET /v1/domains) to get the IDs for your workload domains.
clusterId (optional, VCF 4.4.1.1 and later only)	ID of the specific cluster within a workload domain to apply the custom ISO. If you do not specify a clusterId , the custom ISO will be applied to all clusters in the workload domain. Use the VMware Cloud Foundation API (GET /v1/clusters) to get the IDs for your clusters.
customIsoAbsolutePath	Path to the custom ISO file on the SDDC Manager appliance. For example, /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/VMware-VMvisor-Installer-7.0.0.update01-17325551.x86_64-DellEMC_Customized-A01.iso

Here is an example of a completed JSON template.

```

{
  "esxCustomImageSpecList": [{
    "bundleId": "8c0de63d-b522-4db8-be6c-f1e0ab7ef554",
    "targetEsxVersion": "8.0.1-xxxxxxx",
    "useVcfBundle": false,
    "customIsoAbsolutePath":
"/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-binaries/VMware-VMvisor-
Installer-8.0.0.update01-xxxxxxx.x86_64-DellEMC_Customized-A01.iso"
  }]
}

```

Here is an example of a completed JSON template with multiple ISOs using a single workload domain and specified clusters (VCF 4.4.1.1 and later only).

```

{
  "esxCustomImageSpecList": [
    {
      "bundleId": "aa7b16b1-d719-44b7-9ced-51bb02ca84f4",
      "targetEsxVersion": "8.0.2-xxxxxxx",
      "useVcfBundle": false,
      "domainId": "1b7b16b1-d719-44b7-9ced-51bb02ca84b2",
      "clusterId": "c37b16b1-d719-44b7-9ced-51bb02ca84f4",
      "customIsoAbsolutePath": "/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-
binaries/VMware-ESXi-7.0.2-17867351-DELL.zip"
    },
    {
      "bundleId": "aa7b16b1-d719-44b7-9ced-51bb02ca84f4",
      "targetEsxVersion": "7.0.1-18150133",
      "useVcfBundle": false,
      "domainId": "1b7b16b1-d719-44b7-9ced-51bb02ca84b2",
      "customIsoAbsolutePath": "/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-
binaries/VMware-ESXi-7.0.2-17867351-HP.zip"
    }
  ]
}

```

8. Save the JSON file as `esx-custom-image-upgrade-spec.json` in the `/nfs/vmware/vcf/nfs-mount`.

NOTE

If the JSON file is not saved in the correct directory, the stock VMware Cloud Foundation ISO is used for the upgrade and the custom VIBs are overwritten.

- Set the correct permissions on the `/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json` file:

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

- Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
- In the `lcm.esx.upgrade.custom.image.spec=` parameter, add the path to the JSON file.
For example, `lcm.esx.upgrade.custom.image.spec=/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json`
- In the navigation pane, click **Inventory** > **Workload Domains**.
- On the Workload Domains page, click the domain you are upgrading and then click the **Updates/Patches** tab.
- Schedule the ESXi upgrade bundle.
- Monitor the upgrade progress. See [monitor-update.dita](#).
- After the upgrade is complete, confirm the ESXi version by clicking **Current Versions**. The ESXi hosts table displays the current ESXi version.

Upgrade ESXi with VMware Cloud Foundation Stock ISO and Async Drivers

For clusters in workload domains with vLCM baselines, you can apply the stock ESXi upgrade bundle with specified async drivers.

Download the appropriate async drivers for your hardware on a computer with internet access.

- Download the VMware Cloud Foundation ESXi upgrade bundle. See [downloading-vmware-cloud-foundation-bundles.dita](#).
- Using SSH, log in to the SDDC Manager appliance.
- Create a directory for the vendor provided async drivers under the `/nfs/vmware/vcf/nfs-mount` directory. For example, `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers`.
- Copy the async drivers to the directory you created on the SDDC Manager appliance. For example, you can copy the drivers to the `/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers` directory.
- Change permissions on the directory where you copied the drivers. For example,

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers
```

- Change owner to vcf.
`chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers`
- Create an ESX custom image JSON using the following template.

```
{
  "esxCustomImageSpecList": [{
    "bundleId": "bundle ID of the ESXi bundle you downloaded",
    "targetEsxVersion": "ESXi version for the target VMware Cloud Foundation version",
    "useVcfBundle": true,
```

```
"esxPatchesAbsolutePaths": ["Path_to_Drivers"]
}]
}
```

where

Parameter	Description and Example Value
bundleId	ID of the ESXi upgrade bundle you downloaded. You can retrieve the bundle ID by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the bundle ID. For example, <code>8c0de63d-b522-4db8-be6c-f1e0ab7ef554</code> .
targetEsxVersion	Version of the ESXi upgrade bundle you downloaded. You can retrieve the ESXi target version by navigating to the Lifecycle Management > Bundle Management page and clicking View Details to view the "Update to Version".
useVcfBundle	Specifies whether the ESXi bundle is to be used for the upgrade. Set this to true .
esxPatchesAbsolutePaths	Path to the async drivers on the SDDC Manager appliance. For example, <code>/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/drivers/VMW-ESX-6.7.0-smartpqi-1.0.2.1038-offline_bundle-8984687.zip</code>

Here is an example of a completed JSON template.

```
{
"esxCustomImageSpecList": [{
"bundleId": "411bea6a-b26c-4a15-9443-03f453c68752-apTool",
"targetEsxVersion": "7.0.3-21053776",
"useVcfBundle": true,
"esxPatchesAbsolutePaths": ["/nfs/vmware/vcf/nfs-mount/esx-upgrade-partner-drivers/
drivers/HPE-703.0.0.10.9.5.14-Aug2022-Synergy-Addon-depot.zip"]
}]
}
```

8. Save the JSON file as `esx-custom-image-upgrade-spec.json` in the `/nfs/vmware/vcf/nfs-mount`.

NOTE

If the JSON file is not saved in the correct directory, the stock VMware Cloud Foundation ISO is used for the upgrade and the custom VIBs are overwritten.

9. Set the correct permissions on the `/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json` file:

```
chmod -R 775 /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json
```

10. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.

11. In the `lcm.esx.upgrade.custom.image.spec=` parameter, add the path to the JSON file.
For example, `lcm.esx.upgrade.custom.image.spec=/nfs/vmware/vcf/nfs-mount/esx-custom-image-upgrade-spec.json`
12. In the navigation pane, click **Inventory** > **Workload Domains**.
13. On the Workload Domain page, click the management domain.
14. On the Domain Summary page, click the **Updates/Patches** tab.
15. In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update bundle for VMware ESXi.
16. Monitor the upgrade progress. See [monitor-update.dita](#).
17. After the upgrade is complete, confirm the ESXi version by clicking **Current Versions**. The ESXi hosts table displays the current ESXi version.

Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation 5.2

Prior to VMware Cloud Foundation 5.2.1, workload domains can use either vSphere Lifecycle Manager baselines or vSphere Lifecycle Manager images for ESXi host upgrade. The following procedure describes upgrading ESXi hosts in workload domains that use vSphere Lifecycle Manager images when your target version is VMware Cloud Foundation 5.2.

- Validate that the ESXi passwords are valid.
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.
- You must upgrade NSX and vCenter Server before you can upgrade ESXi hosts with a vSphere Lifecycle Manager image.
- If you want to add firmware to the vSphere Lifecycle Manager image, you must install the Hardware Support Manager from your vendor. See [adding-firmware-and-components-to-a-cluster-image.dita](#).

For information about upgrading ESXi in workload domains that use vSphere Lifecycle Manager baselines when your target version is VMware Cloud Foundation 5.2, see [upgrade-esxi-for-vsphere-lifecycle-manager-baseline-clusters.dita](#). VMware Cloud Foundation 5.2.1 and later support workload domains that include vSphere Lifecycle Manager baseline clusters and vSphere Lifecycle Manager image clusters. If you are upgrading to VMware Cloud Foundation 5.2.1, see [../lifecycle/topics/upgrade-esxi-for-vmware-cloud-foundation-5-2-1.dita](#).

You create a vSphere Lifecycle Manager image for upgrading ESXi hosts using the vSphere Client. During the creation of the image, you define the ESXi version and can optionally add vendor add-ons, components, and firmware. After you extract the vSphere Lifecycle Manager image into SDDC Manager, the ESXi update will be available for the relevant VI workload domains.

1. Log in to the management domain vCenter Server using the vSphere Client.
2. Create a vSphere Lifecycle Manager image.
 - a) Right-click the management domain data center and select **New Cluster**.
 - b) Enter a name for the cluster (for example, `ESXi upgrade image`) and click **Next**.
Keep the default settings for everything except the cluster name.
 - c) Define the vSphere Lifecycle manager image and click **Next**.

Image Element	Description
ESXi Version	From the ESXi Version drop-down menu, select the ESXi version specified in the VMware Cloud Foundation BOM. If the ESXi version does not appear in the drop-down menu, see Working With the vSphere Lifecycle Manager Depot .
Vendor Add-On (optional)	To add a vendor add-on to the image, click Select and select a vendor add-on.

You can customize the image components, firmware, and drivers later.

- d) Click **Finish**.
- e) After the cluster is created successfully, click the **Updates** tab for the new cluster to further customize it, if needed.
- f) Click **Hosts > Image** and then click **Edit**.
- g) Edit the vSphere Lifecycle manager image properties and click **Save**.

You already specified the ESXi version and optional vendor add-on, but you can modify those settings as required.

Image Element	Description
ESXi Version	From the ESXi Version drop-down menu, select the ESXi version specified in the VMware Cloud Foundation BOM. If the ESXi version does not appear in the drop-down menu, see Synchronize the vSphere Lifecycle Manager Depot and Import Updates to the vSphere Lifecycle Manager Depot .
Vendor Add-On (optional)	To add a vendor add-on to the image, click Select and select a vendor add-on.
Firmware and Drivers Add-On (optional)	To add a firmware add-on to the image, click Select . In the Select Firmware and Drivers Addon dialog box, specify a hardware support manager and select a firmware add-on to add to the image. Selecting a firmware add-on for a family of vendor servers is possible only if the respective vendor-provided hardware support manager is registered as an extension to the vCenter Server where vSphere Lifecycle Manager runs.
Components	To add components to the image: <ul style="list-style-type: none"> • Click Show details. • Click Add Components. • Select the components and their corresponding versions to add to the image.

vSphere saves the cluster image.

3. Extract the vSphere Lifecycle Manager image into SDDC Manager.
 - a) In the SDDC Manager UI, click **Lifecycle Management > Image Management** .
 - b) Click **Import Image**.
 - c) In the Option 1 section, select the management domain from the drop-down menu.
 - d) In the Cluster drop-down, select the cluster from which you want to extract the vSphere Lifecycle manager image. For example, *ESXi upgrade image*.
 - e) Enter a name for the cluster image and click **Extract Cluster Image**.

You can view status in the **Tasks** panel.

4. Upgrade ESXi hosts with the vSphere Lifecycle Manager image.
 - a) Navigate to the **Updates** tab of the VI workload domain.
 - b) In the Available Updates section, click **Configure Update**.
 - c) Click **Next**.
 - d) Select the clusters to upgrade and click **Next**.
 The default setting is to upgrade all clusters. To upgrade specific clusters, click **Enable cluster-level selection** and select the clusters to upgrade.
 - e) Select the cluster and the cluster image, and click **Apply Image**.
 - f) Click **Next**.
 - g) Select the upgrade options and click **Next**.

By default, the selected clusters are upgraded in parallel. If you selected more than five clusters to be upgraded, the first five are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Select **Enable Quick Boot** to reduce the upgrade time by skipping the physical reboot of the host.

Select **Enforce Live Patch** when the cluster image includes a Live Patch. With the Enforce Live Patch option, vSphere Lifecycle Manager does not place the hosts in the cluster into maintenance mode, hosts are not rebooted, and there is no need to migrate the virtual machines running on the hosts in the cluster.

Select **Migrate Powered Off and Suspended VMs** to migrate the suspended and powered off virtual machines from the hosts that must enter maintenance mode to other hosts in the cluster.

- h) Review the settings, and click **Finish**.
 VMware Cloud Foundation runs a cluster image hardware compatibility and compliance check. Resolve any reported issues before proceeding.
- i) Click **Schedule Update** and click **Next**.
- j) Select **Upgrade Now** or **Schedule Update** and click **Finish**.
- k) Monitor the upgrade progress. See [monitor-update.dita](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [upgrade-vsan-on-disk-format-versions.dita](#).

Firmware Updates

You can use vSphere Lifecycle Manager images to perform firmware updates on the ESXi hosts in a cluster. Using a vSphere Lifecycle Manager image simplifies the host update operation. With a single operation, you update both the software and the firmware on the host.

To apply firmware updates to hosts in a cluster, you must deploy and configure a vendor provided software module called hardware support manager. The deployment method and the management of a hardware support manager is determined by the respective OEM. For example, the hardware support manager that Dell EMC provides is part of their host management solution, OpenManage Integration for VMware vCenter (OMIVV), which you deploy as an appliance. See [Deploying Hardware Support Managers](#).

You must deploy the hardware support manager appliance on a host with sufficient disk space. After you deploy the appliance, you must power on the appliance virtual machine, log in to the appliance as an administrator, and register the appliance as a vCenter Server extension. Each hardware support manager has its own mechanism of managing firmware packages and making firmware add-ons available for you to choose.

For detailed information about deploying, configuring, and managing hardware support managers, refer to the vendor-provided documentation.

Update License Keys for a Workload Domain

If upgrading from a VMware Cloud Foundation version prior to 5.0, you need to update your license keys to support vSAN 8.x and vSphere 8.x.

You need a new license key for vSAN 8.x and vSphere 8.x. Prior to VMware Cloud Foundation 5.1.1, you must add and update the component license key for each upgraded component in the SDDC Manager UI as described below.

With VMware Cloud Foundation 5.1.1 and later, you can add a component license key as described below, or add a solution license key in the vSphere Client. See [Managing vSphere Licenses](#) for information about using a solution license key for VMware ESXi and vCenter Server. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

You first add the new component license key to SDDC Manager. This must be done once per license instance. You then apply the license key to the component on a per workload domain basis.

1. Add a new component license key to the SDDC Manager inventory.
 - a) In the navigation pane, click **Administration** > **Licensing**.
 - b) On the **Licensing** page, click **+ License Key**.
 - c) Select a product from the drop-down menu.
 - d) Enter the license key.
 - e) Enter a description for the license key.
 - f) Click **Add**.
 - g) Repeat for each license key to be added.
2. Update a license key for a workload domain component.
 - a) In the navigation pane, click **Inventory** > **Workload Domains**.
 - b) On the **Workload Domains** page, click the domain you are upgrading.
 - c) On the **Summary** tab, expand the red error banner, and click **Update Licenses**.
 - d) On the **Update Licenses** page, click **Next**.
 - e) Select the products to update and click **Next**.
 - f) For each product, select a new license key from the list, and select the entity to which the licensekey should be applied and click **Next**.
 - g) On the Review pane, review each license key and click **Submit**.
The new license keys will be applied to the workload domain. Monitor the task in the **Tasks** pane in SDDC Manager.

Upgrade vSphere Distributed Switch versions

[Optional] Upgrade the distributed switch to take advantage of features that are available only in the later versions.

ESXi and vCenter Upgrades are completed.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Upgrade › Upgrade Distributed Switch**
3. Select the vSphere Distributed Switch version that you want to upgrade the switch to and click **Next**

The vSphere Distributed Switch is successfully upgraded.

Upgrade vSAN on-disk format versions

[Optional] Upgrade the vSAN on-disk format version to take advantage of features that are available only in the later versions.

- ESXi and vCenter Upgrades are completed
 - Verify that the disks are in a healthy state. Navigate to the Disk Management page to verify the object status.
 - Verify that your hosts are not in maintenance mode. When upgrading the disk format, do not place the hosts in maintenance mode.
 - Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster. For information about vSAN resynchronization, see vSphere Monitoring and Performance
 - The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure.
1. Navigate to the vSAN cluster.
 2. Click the **Configure** tab.
 3. Under **vSAN**, select **Disk Management**.
 4. Click **Pre-check Upgrade**. The upgrade pre-check analyzes the cluster to uncover any issues that might prevent a successful upgrade. Some of the items checked are host status, disk status, network status, and object status. Upgrade issues are displayed in the **Disk pre-check status text** box.
 5. Click **Upgrade**.
 6. Click **Yes** on the Upgrade dialog box to perform the upgrade of the on-disk format.

vSAN successfully upgrades the on-disk format. The On-disk Format Version column displays the disk format version of storage devices in the cluster

Post Upgrade Steps for NFS-Based VI Workload Domains

After upgrading VI workload domains that use NFS storage, you must add a static route for hosts to access NFS storage over the NFS gateway. This process must be completed before expanding the workload domain.

1. Identify the IP address of the NFS server for the VI workload domain.
2. Identify the network pool associated with the hosts in the cluster and the NFS gateway for the network pool.
 - a) Log in to SDDC Manager.
 - b) Click **Inventory › Workload Domains** and then click the VI workload domain.
 - c) Click the **Clusters** tab and then click an NFS-based cluster.
 - d) Click the **Hosts** tab and note down the network pool for the hosts.
 - e) Click the Info icon next to the network pool name and note down the NFS gateway.

3. Ensure that the NFS server is reachable from the NFS gateway. If a gateway does not exist, create it.
4. Identify the vmknic on each host in the cluster that is configured for NFS traffic.
5. Configure a static route on each host to reach the NFS server from the NFS gateway.

```
esxcli network ip route ipv4 add -g NFS-gateway-IP -n NFS-gateway
```

6. Verify that the new route is added to the host using the NFS vmknic.

```
esxcli network ip route ipv4 list
```

7. Ensure that the hosts in the NFS cluster can reach the NFS gateway through the NFS vmkernel.

For example:

```
vmkping -4 -I vmk2 -s 1470 -d -W 5 10.0.22.250
```

8. Repeat steps 2 through 7 for each cluster using NFS storage.

Independent SDDC Manager Upgrade using the SDDC Manager UI

Once SDDC Manager is upgraded to 5.2 or later, new functionality is introduced that allows you to get the latest SDDC Manager features and security fixes without having to upgrade the entire VMware Cloud Foundation BOM. An independent SDDC Manager release includes a fourth digit in its version number, for example SDDC Manager 5.2.0.1.

- Download the SDDC Manager bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- SDDC Manager must be version 5.2 or later.

You can upgrade SDDC Manager without upgrading the full VCF BOM when:

- The target version of SDDC Manager is compatible with all the BOM product versions running in your current environment (management and workload domains).
- There is a supported upgrade path from your current SDDC Manager version to the target SDDC Manager version.

NOTE

You can use the SDDC Manager upgrade functionality to upgrade SDDC Manager even when the target version of SDDC Manager is part of a full VCF BOM release, as long as it is compatible.

Updating SDDC Manager without upgrading the full VCF BOM, does not change the version of the management domain.

1. In the navigation pane, browse to **Lifecycle Management** > **SDDC Manager**.
The UI displays available SDDC Manager updates that are either SDDC Manager only updates or SDDC Manager updates that are part of a full VCF BOM update.
2. Review and address any compatibility warnings.
3. Click **Run Precheck**.
Resolve any precheck issues before proceeding.
4. Schedule the update to run now or at a specific time and click **Start Update**.
When the update completes successfully, you are logged out of the SDDC Manager UI and must log in again.

Flexible BOM Upgrade in VMware Cloud Foundation

Once SDDC Manager is upgraded to version 5.2 or later, new functionality is introduced to the upgrade planner that allows you to select specific target versions for each VMware Cloud Foundation component you want to upgrade.

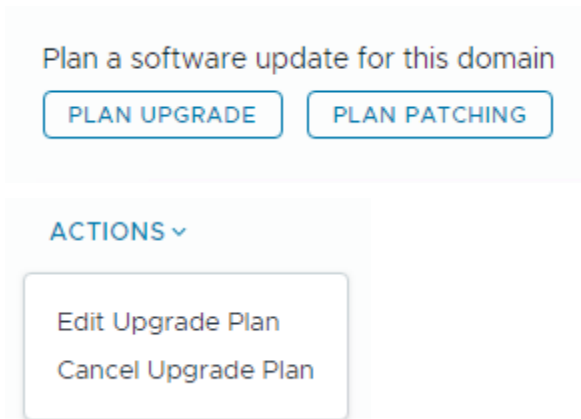
- Download the bundles for the target versions of each VCF component. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- SDDC Manager must be version 5.2 or later.

You can use the upgrade planner to select any supported version for each of the VMware Cloud Foundation BOM components. This includes async patch versions as well as VCF BOM versions. To plan an upgrade when SDDC Manager does not have internet access, see [Offline Download of Flexible BOM Upgrade Bundles](#).

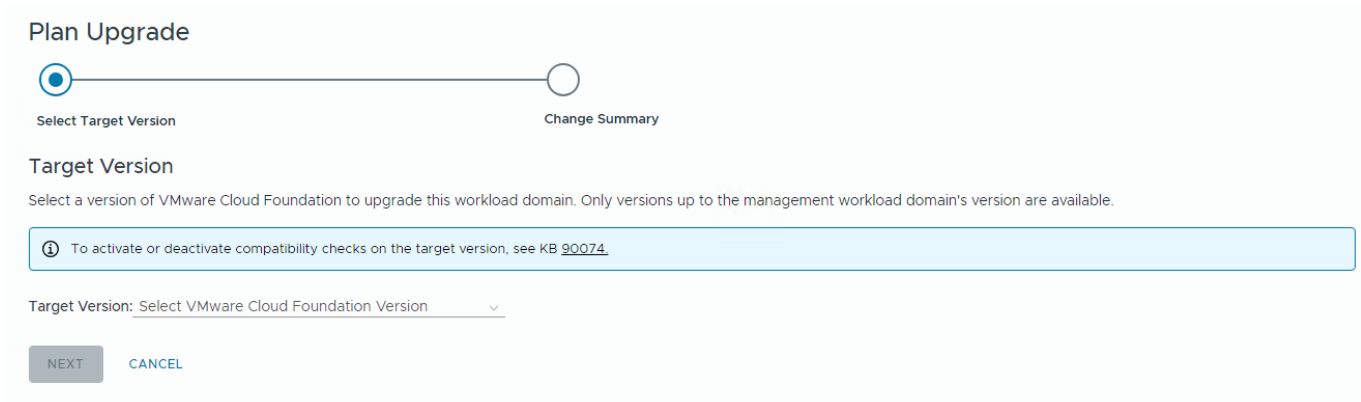
1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with an upgrade.

4. In the Available Updates section, click **Plan Upgrade** create a new upgrade plan or select **Edit Upgrade Plan** from the **Actions** menu to modify an upgrade plan.



5. Select the target version of VMware Cloud Foundation from the drop-down menu and click **Next**.



6. Click **Customize Upgrade** to select specific target versions for each VCF BOM component.
7. Use the drop-down menus in the Target Version column to select a target version for each component and then click **Validate Selection**.

Change Summary


The following product upgrades will be planned out. The target versions of individual products can be customized if necessary.

Software Component	Current Version	Current Build	Target Version	Target Build
VMware ESXi	7.0.3	20328353	8.0.3-24017672 ▾	24017672
VMware NSX	3.2.1.2.0	20541212	4.2.0.0.0-24013905 ▾	24013905
VMware vCenter Server Appliance	7.0.3.01000	20395099	8.0.3.00000-24017671 ▾	24017671

CONFIRM **BACK** **CANCEL**

8. After validation succeeds, click **Confirm**.
9. Review the update sequence based on your target version selections and click **Done**.


Plan Upgrade


 Upgrade plan confirmed. See your upgrade summary below.




Selected Target VMware Cloud Foundation version
5.2.0.0

Update sequence preview

- Step 1  **NSX_T_MANAGER Update Bundle 4.2.0.0.0**

The upgrade bundle for VMware NSX Data Center 4.2.0.0. Customers are strongly encouraged to run the NSX Upgrade Evaluation Tool. For more information, see <https://docs.vmware.com/en/VMware-NSX/4.2/rn/vmware-nsxt-data-center-42-release-notes/index.html>.
- Step 2  **VMware vCenter Server Update Bundle 8.0.3.00000**

The upgrade bundle for VMware vCenter Server 8.0.3. For more information, see <https://docs.vmware.com/en/VMware-vSphere/8.0.3/rn/vsphere-vcserver-803-release-notes.html>.
- Step 3  **VMware ESXi Server Update Bundle 8.0.3**

The upgrade bundle for VMware ESXi 8.0.3. For more information, see <https://docs.vmware.com/en/VMware-vSphere/8.0.3/rn/vsphere-esxi-803-release-notes.html>.

DONE **BACK AND EDIT UPGRADE PLAN**

10. In the Available Updates screen, click **Schedule Update** or **Update Now** to update the first component. Continue to update the VCF BOM components until they are all updated.

NOTE

If SDDC Manager does not have internet access, you need to perform additional steps before you can start updating. See [Offline Download of Flexible BOM Upgrade Bundles](#).

Patching the Management and Workload Domains

Once SDDC Manager is upgraded to 5.2 or later, a new option for patching VMware Cloud Foundation components is available in the SDDC Manager UI.



- Download the async patch bundles. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- SDDC Manager must be version 5.2 or later. See [Apply the VMware Cloud Foundation 5.2.x Upgrade Bundle](#).

The patch planner provides the ability to apply async patches to workload domain components. If you are connected to the online depot, async patches are available in the patch planner. If you do not have access to the online depot, use the Bundle Transfer Utility to download async patches and add them to an offline depot or upload them directly to SDDC Manager.

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the domain you are patching and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with an upgrade.

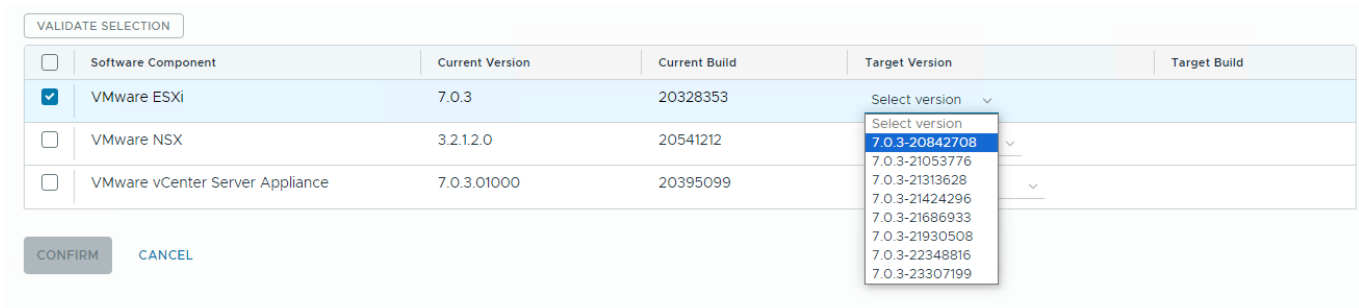
4. In the Available Updates section, click **Plan Patching** create a new patching plan or select **Edit Patching Plan** from the **Actions** menu to modify a patching plan.



NOTE

You cannot plan patching if you have an existing upgrade plan. Cancel the upgrade plan to create a patching plan.

5. Select the components to patch and the target versions and then click **Validate Selection**.



NOTE

When you select a target vCenter version, the UI indicates which versions support vCenter Reduced Downtime Upgrade (RDU).

6. After validation succeeds, click **Confirm**.
7. Review the update sequence based on your target version selections and click **Done**.

Plan Patching

Update sequence preview

Step 1 ESX_HOST 7.0.3

This VMware Software Upgrade bundle contains VMware ESXi 7.0 Update 3i. For more information, see <https://docs.vmware.com/en/VMware-vSphere/7.0/rn/vsphere-esxi-70u3i-release-notes.html>

DONE BACK AND EDIT PATCHES

8. In the Available Updates screen, click **Schedule Update** or **Update Now** to update the first component. Continue to update the VCF BOM components until they are all updated.

Troubleshooting for Upgrading VMware Cloud Foundation

A library of troubleshooting processes that may be referenced during the VMware Cloud Foundation upgrade as appropriate.

SDDC Manager Troubleshooting

A library of SDDC Manager troubleshooting processes that may be referenced during upgrade as appropriate.

On-demand pre-checks for vCenter bundle might fail

The bundle pre-check failure can occur in a specific scenario. When SDDC Manager is upgraded to VMware Cloud Foundation 5.0.0.x from 4.5.x, and BOM components are not upgraded to VMware Cloud Foundation 5.0.0.x and Customer downloads the bundles for VMware Cloud Foundation 5.1.0.0 and runs the pre-check by selecting target version as 5.1.0.0.

The format of the vCenter Server bundle is modified starting from VMware Cloud Foundation 5.1. The new bundle is a unified bundle that bundles both the .iso and .zip files for the Target vCenter Server build. This unified bundle can be used for both major and minor vCenter Server upgrades. The SDDC Manager needs to be at least at the 5.1 version to understand the new format and run the prechecks. As VMware Cloud Foundation 5.0.0.0 does not understand the format, the bundle pre-check will fail.

Error Message: Upgrade Bundle Validation

1. Upgrade the SDDC Manager to VMware Cloud Foundation 5.1.0.0 and run the on-demand prechecks for vCenter Server in VMware Cloud Foundation 5.1.0.0.

<https://kb.vmware.com/s/article/94862>

SDDC Manager bundle pre-check failure when upgrading to VMware Cloud Foundation 5.1

SDDC Manager Pre-check fails

SDDC Manager Pre-check "Upgrade Bundle Download Status" fails with an error

- "Could not find bundle for SDDC_MANAGER upgrade to version 5.1.0.0-<build_number>".

From VMware Cloud Foundation 5.1 onwards, we are deprecating the Config Drift bundle. However, the previously released VCF versions expect that a config drift bundle will be applied as part of a target release and hence indicate this as a pre-check failure.

This pre-check failure can be ignored for VCF 5.1+, and it is safe to proceed with the upgrade despite this bundle pre-check failure.

Extra RPM packages on SDDC Manager may cause upgrade failure

SDDC Manager upgrade may fail if some RPMs on the current SDDC Manager are incompatible with those on the upgraded SDDC Manager. In `/var/log/vmware/capengine/cap-update/install-*`,

You may see a message like:

- package `systemd-udev-247.13-4.ph4.x86_64` requires `libsystemd-shared-247.so()(64bit)`, but none of the providers can be installed.
- package `systemd-247.13-4.ph4.x86_64` requires `libcrypto.so.3()(64bit)`, but none of the providers can be installed.
- package `rpm-4.16.1.3-17.ph4.x86_64` requires `libcrypto.so.3()(64bit)`, but none of the providers can be installed

RPMs may have been left behind by previous upgrades or greenfield deployments, or a user has implicitly or explicitly installed RPMs that prevent the upgrade

The workaround is to uninstall RPMs that are causing this upgrade conflict manually.

False warning for missing compatibility data in plan upgrade wizard

When no compatibility data is missing, an incorrect warning message is populated

A warning message with an empty product list in the plan upgrade wizard appears

- "Unable to verify the compatibility for the following product versions. Please check the product documentation before proceeding to upgrade:"

Users can ignore the warning and is not blocked.

Updating licenses for a WLD shows insufficient license error

When the 'Update Licenses' operation is performed for a Workload Domain, in certain cases, the incorrect quantity of licenses is shown in the 'Available quantity' field

This is due to a miscalculation in the no. of available licenses. Along with the incorrect quantity, an error alert might also be displayed saying,

- 'License key has insufficient license.'

A miscalculation in the code for the number of available licenses is causing the error alert to appear.

The users can simply choose to ignore the incorrect license count in the 'Available quantity' field when assigning the license. Also, the error alert should be ignored as it does not prohibit the user from moving forward. Users can proceed with the addition of a license even with the error alert. If there are sufficient licenses available, the operation will succeed.

vCenter Troubleshooting

A library of vCenter troubleshooting processes that may be referenced during upgrade as appropriate.

vCenter Server Upgrade Failed Due to Reuse of Temporary IP Address

vCenter Server Upgrade Failed Due to Reuse of Temporary IP Address with error "Cannot run the revert networking command. revert_networking.py doesn't exist on target VC" or "VC upgrade is failing during Install-"target vc upgrade precheck stage failing"

Reuse of temporary IP address causes an arp cache issue. Reset the arp cache on the management domain vCenter Server.

Customers who have fewer Temporary IP Addresses than vCenter Servers that are conducting a parallel upgrade have the highest likelihood of impact.

1. SSH to the management domain vCenter Server as root.
2. Run the following

```
ip -s -s neigh flush all
```

Async Patch Tool

The Async Patch Tool is a utility that allows you to apply critical patches to certain VMware Cloud Foundation components (NSX Manager, vCenter Server, and ESXi) outside of VMware Cloud Foundation releases. The Async Patch Tool also supports VxRail Manager patching of VMware Cloud Foundation on Dell VxRail.

Async Patch Tool 1.2

The Async Patch Tool is a utility that allows you to apply critical patches to certain VMware Cloud Foundation components (NSX Manager, vCenter Server, and ESXi) outside of VMware Cloud Foundation releases. The Async Patch Tool also supports VxRail Manager patching of VMware Cloud Foundation on Dell VxRail.

For example, you could use the Async Patch Tool to get a vCenter Server patch that addresses a critical security issue as described in a VMware Security Advisory (VMSA). You use the Async Patch Tool to download the patch and upload it to the internal LCM repository on the SDDC Manager appliance. Then you use the SDDC Manager UI to apply the patch.

NOTE

Patched components will have different versions than those listed in the Bill of Materials (BOM).

The Async Patch Tool is supported with VMware Cloud Foundation 4.2.1 and later. This release also supports VxRail Manager patching of VMware Cloud Foundation on VxRail.

NOTE

Standalone ESXi async patches are not supported with VMware Cloud Foundation on Dell VxRail. Applying a VxRail Manager async patch also patches ESXi.

The process for downloading and uploading patches varies depending on whether or not the SDDC Manager appliance has access to the internet. If the SDDC Manager appliance has access to the internet (online), you perform all Async Patch Tool operations from the SDDC Manager appliance. If the SDDC Manager appliance does not have access to the internet (offline), you perform some operations from a computer with internet access and some operations from the SDDC Manager appliance. See:

- [Apply an Async Patch to VMware Cloud Foundation in Online Mode](#)
- [Apply an Async Patch to VMware Cloud Foundation in Offline Mode](#)

IMPORTANT

SDDC Manager 5.2 and later support applying async patches directly from the SDDC Manager UI. See [Patching the Management and Workload Domains](#).

After you apply an async patch, and one or more of the VMware Cloud Foundation component versions deviates from the BOM, you may have to use the Async Patch Tool to enable an upgrade to a later version of VMware Cloud Foundation.

Target version	Requires use of Async Patch Tool?
VMware Cloud Foundation 4.y	<p>Yes.</p> <p>If you are upgrading an async patched system from VMware Cloud Foundation 4.x to 4.y, you must use the Async Patch Tool to enable the upgrade.</p> <p>For example, if you apply a vCenter Server patch to a VMware Cloud Foundation 4.5.0 instance, you must use the Async Patch Tool to enable an upgrade to VMware Cloud Foundation 4.5.2. See:</p>

Table continued on next page

Continued from previous page

Target version	Requires use of Async Patch Tool?
	<ul style="list-style-type: none"> • Upgrade an Async Patched Version of VMware Cloud Foundation in Online Mode • Upgrade an Async Patched Version of VMware Cloud Foundation in Offline Mode
VMware Cloud Foundation 5.x	<p>No.</p> <p>If you are upgrading an async patched system from VMware Cloud Foundation 4.x to 5.x or 5.x to 5.x, you do not need to use the Async Patch Tool to enable the upgrade. Upgrades to 5.x are automatically enabled and you can upgrade using the SDDC Manager UI or the Bundle Transfer Utility. See the <i>VMware Cloud Foundation Lifecycle Management Guide</i>.</p> <p>NOTE You should still use the Async Patch Tool to deactivate all async patches and run an inventory sync before upgrading to VMware Cloud Foundation 5.x. See "VCF Async Patch Tool Options" in the Async Patch Tool documentation for more information.</p>

For information about known issues, see the *Async Patch Tool Release Notes*.

Apply an Async Patch to VMware Cloud Foundation in Online Mode

If your SDDC Manager appliance has a connection to the internet (either directly or through a proxy server), you can run the Async Patch Tool from the SDDC Manager appliance to download and enable an async patch. Once the patch is successfully enabled, you can use the SDDC Manager UI to apply the patch to all workload domains.

- Refer to [KB 88287](#) to ensure that the async patch is supported with your version of VMware Cloud Foundation. Contact VMware Support if you have questions about the available async patches and which versions of VMware Cloud Foundation support them.
- You must have the latest version of the Async Patch Tool.

NOTE

If an existing or older version of the Async Patch Tool exists in the directory, you will need to remove these files before downloading the latest version of the Async Patch Tool.

```
rm -r /home/vcf/asyncPatchTool
```

```
rm -r <outputdirectory>
```

The default directory is `/home/vcf/apToolBundles` if `outputDirectory` was not specified when the Async Patch Tool was previously run.

- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Async Patch Tool for long-running operations.
- The Async Patch Tool is supported with VMware Cloud Foundation 4.2.1 and later. This release also supports ESXi and VxRail Manager patching of VMware Cloud Foundation on VxRail.

1. Download the most recent version of the Async Patch Tool to a computer that has access to the SDDC Manager appliance.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.

- b) Click your current version of VMware Cloud Foundation.
 - c) Click **Drivers & Tools**.
 - d) Click the download icon for the Async Patch Tool.
2. Copy the Async Patch Tool to the SDDC Manager appliance and configure it for use.

- a) SSH in to the SDDC Manager appliance using the `vcf` user account.
- b) Create the `asyncPatchTool` directory.

```
mkdir /home/vcf/asyncPatchTool
```

- c) Copy the Async Patch Tool file (`vcf-async-patch-tool-<version>.tar.gz`) that you downloaded in step 1 to the `/home/vcf/asyncPatchTool` directory.
- d) Navigate to `/home/vcf/asyncPatchTool` and extract the contents of `vcf-async-patch-tool-<version>.tar.gz`.

```
tar -xvf vcf-async-patch-tool-<version>.tar.gz
```

- e) Set the permissions for the `asyncPatchTool` directory.

```
cd /home/vcf/
chmod -R 755 asyncPatchTool
chown -R vcf:vcf asyncPatchTool
```

3. List the available async patches.

- a) Navigate to `/home/vcf/asyncPatchTool/bin`.
- b) Run the following command:

```
./vcf-async-patch-tool --listAsyncPatch --du broadcom_support_email
```

Replace *customer_connect_email* with your Broadcom Support portal email address.

Optionally, you can use the `--sku` and `--productType` options to filter the list of patches. See [VCF Async Patch Tool Options](#) for details.

`--outputDirectory` is optional and can be used to specify a location for the download. Select a directory that has enough free space for the bundle. If you do not specify a location, the Async Patch Tool displays the default location in its output. For example: `/root/apToolBundles`.

NOTE

If you connect to the internet through a proxy server, use the `--proxyServer`, `--ps` option to specify the FQDN and port of the proxy server. For example, `--proxyServer FQDN:port`.

- c) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
- d) Enter `Y` or `N` to choose whether or not to participate in the Customer Experience Improvement Program (CEIP).
- e) Enter your Broadcom Support portal password.

The Async Patch Tool lists all available async patches.

4. (VxRail async patch only) Copy the VxRail async patch-specific partner bundle metadata file using [KB 91830](#).
5. Enable an async patch.

- a) Run the following command:

```
VMware Cloud Foundation:
```

```
./vcf-async-patch-tool -e --patch product:version --du broadcom_support_email
--sddcSSOUser SSOuser --sddcSSHUser vcf --it ONLINE
```

VMware Cloud Foundation on Dell EMC VxRail:

```
./vcf-async-patch-tool -e --patch product:version --du broadcom_support_email
--pdu dell_emc_depot_email --sddcSSOUser SSOuser --sddcSSHUser vcf --it ONLINE
```

- Replace *product:version* with the product and version of a patch retrieved in step 3. For example: VCENTER:7.0.3.00300-19234570.
- Replace *broadcom_support_email* with your Broadcom Support portal email address.
- Replace *dell_emc_depot_email* with your Dell EMC depot email address. (VxRail only)
- Replace *SSOuser* with the management domain SSO user account, for example, administrator@vsphere.local.

NOTE

If you connect to the internet through a proxy server, use the `--proxyServer`, `--ps` option to specify the FQDN and port of the proxy server. For example, `--proxyServer FQDN:port`.

- b) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
- c) Read the information and enter `Y` to acknowledge the pre-requisites.
- d) Enter `Y` or `N` to choose whether or not to participate in the VMware Customer Experience Improvement Program (CEIP).
- e) Enter the password for the super user (*vcf*) account.
- f) Enter the password for the root user account.
- g) Enter the password for the management domain SSO user account.
- h) Enter your Broadcom Support portal password.
- i) If the product type is **VX_MANAGER**, enter your Dell EMC Depot user name and password. (VxRail only)

The Async Patch Tool downloads the patch and uploads it to the internal LCM repository on the SDDC Manager appliance.

6. Log in to the SDDC Manager UI and apply the async patch to all workload domains.
 - For clusters in workload domains with vSphere Lifecycle Manager baselines, you can upgrade ESXi to the async patch version with a custom ISO from your vendor. See "Upgrade ESXi with Custom ISOs" in *VMware Cloud Foundation Lifecycle Management*.
 - For clusters in workload domains with vSphere Lifecycle Manager images, you can upgrade ESXi to the async patch version by following the procedure "Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation" in *VMware Cloud Foundation Lifecycle Management*.
7. After the async patch is successfully applied, use the Async Patch Tool to deactivate the patch.
 - a) SSH in to the SDDC Manager appliance using the *vcf* user account.
 - b) Navigate to `/home/vcf/asyncPatchTool/bin`.
 - c) Run the following command:


```
./vcf-async-patch-tool --disableAllPatches --sddcSSOUser SSOuser --sddcSSHUser vcf
```

Replace *SSOuser* with the management domain SSO user account, for example, administrator@vsphere.local.
 - d) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
 - e) Enter `Y` or `N` to choose whether or not to participate in the Customer Experience Improvement Program (CEIP).
 - f) Enter the password for the super user (*vcf*) account.

- g) Enter the password for the root user account.
 - h) Enter the password for the management domain SSO user account.
8. (VxRail only) If you applied an async patch to VMware Cloud Foundation on Dell VxRail, reconnect SDDC Manager to the VMware Depot using the SDDC Manager UI or VMware Cloud Foundation API.

Starting with VMware Cloud Foundation 5.2, if you applied a vCenter Server or NSX Manager async patch to the management domain, any new workload domains that you deploy will include the patched version of vCenter Server and/or NSX Manager.

For versions of VMware Cloud Foundation earlier than 5.2, new workload domains will not include async patch versions of vCenter Server or NSX Manager. Use this procedure to apply the async patch(es) to the new workload domain.

NOTE

After you update the hosts in a workload domain to an async patch version of ESXi, any new hosts that you add to the workload domain must use the async patch version of ESXi and not the version listed in the VMware Cloud Foundation BOM.

Apply an Async Patch to VMware Cloud Foundation in Offline Mode

If your SDDC Manager appliance does not have a connection to the internet, you can run the Async Patch Tool from a computer that does. Download an async patch, copy the patch and the Async Patch Tool to the SDDC Manager appliance, and enable the patch. You can then use the SDDC Manager UI to apply the patch to all workload domains.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy server) for downloading the bundles.
- The computer must have Java 8 or Java 11.
- A Windows or Linux computer with access to the SDDC Manager appliance for uploading the bundles.
- Refer to [KB 88287](#) to ensure that the async patch is supported with your version of VMware Cloud Foundation. Contact VMware Support if you have questions about the available async patches and which versions of VMware Cloud Foundation support them.
- You must have the latest version of the Async Patch Tool.

NOTE

If an existing or older version of the Async Patch Tool exists in the directory, you will need to remove these files from both the Linux or Windows computer and the SDDC manager before downloading the latest version of the Async Patch Tool.

```
rm -r <AP Tool directory>
```

```
rm -r <outputdirectory>
```

The default directory is `/home/vcf/apToolBundles` if `outputDirectory` was not specified when the Async Patch Tool was previously run.

- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Async Patch Tool for long-running operations.
 - The Async Patch Tool is supported with VMware Cloud Foundation 4.2.1 and later. This release also supports ESXi and VxRail Manager patching of VMware Cloud Foundation on VxRail.
1. Download the most recent version of the Async Patch Tool to a computer that has access to the internet.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
 - b) Click your current version of VMware Cloud Foundation.
 - c) Click **Drivers & Tools**.
 - d) Click the download icon for the Async Patch Tool.
 2. Extract `vcf-async-patch-tool-<version>.tar.gz`.

3. Navigate to `vcf-async-patch-tool-<version>/bin` and confirm that you have execute permissions.
4. List the available async patches.

- a) Run the following command:

Linux:

```
./vcf-async-patch-tool --listAsyncPatch --du broadcom_support_email
```

Windows:

```
vcf-async-patch-tool.bat --listAsyncPatch --du broadcom_support_email
```

Replace *broadcom_support_email* with your Broadcom Support portal email address.

Optionally, you can use the `--sku` and `--productType` options to filter the list of patches. See [VCF Async Patch Tool Options](#) for details.

`--outputDirectory` is optional and can be used to specify a location for the download. Select a directory that has enough free space for the bundle. If you do not specify a location, the Async Patch Tool displays the default location in its output. For example: `/root/apToolBundles`.

NOTE

If you connect to the internet through a proxy server, use the `--proxyServer`, `--ps` option to specify the FQDN and port of the proxy server. For example, `--proxyServer FQDN:port`.

- b) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
- c) Enter `Y` or `N` to choose whether or not to participate in the Customer Experience Improvement Program (CEIP).
- d) Enter your Broadcom Support portal password.

The Async Patch Tool lists all available async patches.

5. (VxRail async patch only) Copy the VxRail async patch-specific partner bundle metadata file using [KB 91830](#).
6. Download an async patch.

- a) Run the following command:

Linux:

```
./vcf-async-patch-tool -d --patch product:version --du broadcom_support_email
--sku sku_type --pdu dell_emc_depot_email --sddcManagerVersion
current_sddc_version
```

Windows:

```
vcf-async-patch-tool.bat -d --patch product:version --du broadcom_support_email
--sku sku_type --pdu dell_emc_depot_email --sddcManagerVersion
current_sddc_version
```

- Replace *product:version* with the product and version of a patch retrieved in step 4. For example: `VCENTER:7.0.3.00300-19234570`.
- Replace *broadcom_support_email* with your Broadcom Support portal email address.
- Replace *sku_type* with `VCF` or `VCF_ON_VXRAIL`.
- Replace *dell_emc_depot_email* with your Dell EMC Depot email address. (VxRail only)
- Replace *current_sddc_version* with your current version of SDDC Manager. For example: `4.5.0.0`. This is optional, but limits the number of bundles that are downloaded to only those that are applicable to your current version of SDDC Manager.
- `--outputDirectory` is optional and can be used to specify a location for the download. Select a directory that has enough free space for the bundle. If you do not specify a location, the Async Patch Tool displays the default location in its output. For example: `/root/apToolBundles`.

NOTE

If you connect to the internet through a proxy server, use the `--proxyServer`, `--ps` option to specify the FQDN and port of the proxy server. For example, `--proxyServer FQDN:port`.

- b) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
 - c) Enter your Broadcom Support portal password.
 - d) If the product type is **VX_MANAGER**, enter your Dell EMC Depot user name and password. (VxRail only)
- The Async Patch Tool downloads the patch and required artifacts (for example, the LCM manifest).
7. Copy the patch and set permissions.
 - a) Copy the entire output directory (for example, `apToolBundles`) to the SDDC Manager appliance.

You can select any location that has enough free space available, for example, `/nfs/vmware/vcf/nfs-mount/apToolBundles`.
 - b) SSH in to the SDDC Manager appliance using the `vcf` user account.
 - c) Navigate to `/nfs/vmware/vcf/nfs-mount/apToolBundles`.

If you copied the output directory to a different location, navigate to that directory instead.
 - d) Run the following commands:


```
chmod -R 755 apToolBundles
chown -R vcf:vcf apToolBundles
```
 8. Copy the Async Patch Tool to the SDDC Manager appliance and configure it for use.
 - a) SSH in to the SDDC Manager appliance using the `vcf` user account.
 - b) Create the `asyncPatchTool` directory.


```
mkdir /home/vcf/asyncPatchTool
```
 - c) Copy the entire contents of the Async Patch Tool directory from the computer with internet access to the `/home/vcf/asyncPatchTool` directory on the SDDC Manager appliance.
 - d) Set the permissions for the `asyncPatchTool` directory.


```
cd /home/vcf/
chmod -R 755 asyncPatchTool
chown -R vcf:vcf asyncPatchTool
```
 9. Enable an async patch.
 - a) Navigate to `/home/vcf/asyncPatchTool/bin` and run the following command:


```
./vcf-async-patch-tool -e --patch product:version --sddcSSOUser SSOuser --sddcSSHUser vcf --outputDirectory bundleDirectory --it OFFLINE
```

 - Replace `product:version` with the product and version of a patch retrieved in step 4. For example: `VCENTER:7.0.3.00300-19234570`.
 - Replace `SSOuser` with the management domain SSO user account, for example, `administrator@vsphere.local`.
 - Replace `bundleDirectory` with the location of the bundle directory from step 6. For example, `/nfs/vmware/vcf/nfs-mount/apToolBundles`.
 - b) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.

- c) Enter `Y` or `N` to choose whether or not to participate in the Customer Experience Improvement Program (CEIP).
- d) Read the information and enter `Y` to acknowledge the pre-requisites.
- e) Enter the password for the super user (`vcf`) account.
- f) Enter the password for the root user account.
- g) Enter the password for the management domain SSO user account.

The Async Patch Tool uploads the patch to the internal LCM repository on the SDDC Manager appliance.

10. Log in to the SDDC Manager UI and apply the async patch to all workload domains.
 - For clusters in workload domains with vSphere Lifecycle Manager baselines, you can upgrade ESXi to the async patch version with a custom ISO from your vendor. See "Upgrade ESXi with Custom ISOs" in *VMware Cloud Foundation Lifecycle Management*.
 - For clusters in workload domains with vSphere Lifecycle Manager images, you can upgrade ESXi to the async patch version by following the procedure "Upgrade ESXi with vSphere Lifecycle Manager Images for VMware Cloud Foundation" in *VMware Cloud Foundation Lifecycle Management*.
11. After the async patch is successfully applied, use the Async Patch Tool to deactivate the patch.
 - a) SSH in to the SDDC Manager appliance using the `vcf` user account.
 - b) Navigate to `/home/vcf/asyncPatchTool/bin`.
 - c) Run the following command:


```
./vcf-async-patch-tool --disableAllPatches --sddcSSOUser SSOuser --sddcSSHUser vcf
```

Replace `SSOuser` with the management domain SSO user account, for example, `administrator@vsphere.local`.
 - d) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
 - e) Enter `Y` or `N` to choose whether or not to participate in the Customer Experience Improvement Program (CEIP).
 - f) Enter the password for the super user (`vcf`) account.
 - g) Enter the password for the root user account.
 - h) Enter the password for the management domain SSO user account.

Starting with VMware Cloud Foundation 5.2, if you applied a vCenter Server or NSX Manager async patch to the management domain, any new workload domains that you deploy will include the patched version of vCenter Server and/or NSX Manager.

For versions of VMware Cloud Foundation earlier than 5.2, new workload domains will not include async patch versions of vCenter Server or NSX Manager. Use this procedure to apply the async patch(es) to the new workload domain.

NOTE

After you update the hosts in a workload domain to an async patch version of ESXi, any new hosts that you add to the workload domain must use the async patch version of ESXi and not the version listed in the VMware Cloud Foundation BOM.

Upgrade an Async Patched Version of VMware Cloud Foundation in Online Mode

After you apply an async patch, and one or more of the VMware Cloud Foundation component versions deviates from the BOM, you may have to use the Async Patch Tool to enable an upgrade to a later version of VMware Cloud Foundation.

- Refer to [KB 88287](#) to ensure that the target version of VMware Cloud Foundation is supported for upgrade. Contact Broadcom Support if you have questions about support for your upgrade path.
- You must have the latest version of the Async Patch Tool.

NOTE

If an existing or older version of the Async Patch Tool exists in the directory, you will need to remove these files from both the Linux or Windows computer and the SDDC manager before downloading the latest version of the Async Patch Tool.

```
rm -r <AP Tool directory>
```

```
rm -r <outputdirectory>
```

The default directory is `/home/vcf/apToolBundles` if `outputDirectory` was not specified when the Async Patch Tool was previously run.

- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Async Patch Tool for long-running operations.
- The Async Patch Tool is supported with VMware Cloud Foundation 4.2.1 and later. This release also supports ESXi and VxRail Manager patching of VMware Cloud Foundation on VxRail.
- You are upgrading from VMware Cloud Foundation 4.x to VMware Cloud Foundation 4.y.

NOTE

This procedure is not required if you are upgrading from VMware Cloud Foundation 4.x to VMware Cloud Foundation 5.0.

After you apply an async patch, and one or more of the VMware Cloud Foundation component versions deviates from the BOM, the process for upgrading to a later version of VMware Cloud Foundation depends on the target version of VMware Cloud Foundation.

Target Version	Upgrade Process
VMware Cloud Foundation 4.y	<p>You must use the Async Patch Tool to enable an upgrade to a later version of VMware Cloud Foundation.</p> <p>For example, if you apply a vCenter Server patch to a VMware Cloud Foundation 4.3.1 instance, you must use the Async Patch Tool to enable upgrade to VMware Cloud Foundation 4.4.</p> <p>The following procedure describes this process.</p>
VMware Cloud Foundation 5.0	<p>Upgrades to 5.0 are automatically enabled and you can upgrade using the SDDC Manager UI.</p> <p>This process is described in the <i>VMware Cloud Foundation Lifecycle Management Guide</i>.</p> <p>NOTE You should still use the Async Patch Tool to deactivate all async patches and run an inventory sync before upgrading to VMware Cloud Foundation 5.0. See "VCF Async Patch Tool Options" for more information.</p>

If your SDDC Manager appliance has a connection to the internet (either directly or through a proxy server), you can run the Async Patch Tool from the SDDC Manager appliance to download and enable upgrade bundles. Once the bundles are successfully enabled, you can use the SDDC Manager UI to upgrade VMware Cloud Foundation.

If you configured a connection to the VMware Depot in the SDDC Manager UI (**Administration > Repository Settings**), then the upgrade bundles get downloaded automatically. However, you cannot apply the bundles to a workload domain until use this procedure to enable the upgrade.

1. SSH in to the SDDC Manager appliance using the `vcf` user account.
2. Navigate to `/home/vcf/asyncPatchTool/bin`.
3. Enable the upgrade.

- a) Run the following command:

VMware Cloud Foundation:

```
./vcf-async-patch-tool --enableVCFUpgradetargetVcfVersion--dubroadcom_support_email--sddcSSOUsersssoUsername--sddcSSHUser vcf --it ONLINE
```

VMware Cloud Foundation on Dell EMC VxRail:

```
./vcf-async-patch-tool --enableVCFUpgradetargetVcfVersion--dubroadcom_support_email--pdudell_emc_depot_email--sddcSSOUsersssoUsername--sddcSSHUser vcf --it ONLINE
```

- Replace `targetVcfVersion` with the target version of VMware Cloud Foundation. For example: `4.4.0.0`.
- Replace `broadcom_support_email` with your Broadcom Support portal email address.
- Replace `dell_emc_depot_email` with your Dell EMC depot email address. (VxRail only)
- Replace `ssoUsername` with the management domain SSO user account, for example, `administrator@vsphere.local`.
- `--outputDirectory` is optional and can be used to specify a location for the bundle download. If you do not specify a location, the Async Patch Tool download to the default location, `/root/apToolBundles`.

NOTE

Ensure that the output directory has enough free space for the bundles.

NOTE

If you connect to the internet through a proxy server, use the `--proxyServer`, `--ps` option to specify the FQDN and port of the proxy server. For example, `--proxyServer FQDN:port`.

- b) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
- c) Enter `Y` or `N` to choose whether or not to participate in the Customer Experience Improvement Program (CEIP).
- d) Read the information and enter `Y` to acknowledge the pre-requisites.
- e) Enter the password for the super user (`vcf`) account.
- f) Enter the password for the root user account.
- g) Enter the password for the management domain SSO user account.
- h) Enter your Broadcom Support portal password.
- i) If the product type is **VX_MANAGER**, enter your Dell EMC Depot user name and password. (VxRail only)
- j) Enter `Y` or `N` to choose whether or not to download vRealize bundles.

The Async Patch Tool determines which bundles are required, downloads the bundles, and uploads them to the internal LCM repository on the SDDC Manager appliance.

4. Log in to the SDDC Manager UI and apply the bundles to all workload domains.

Upgrade an Async Patched Version of VMware Cloud Foundation in Offline Mode

After you apply an async patch, and one or more of the VMware Cloud Foundation component versions deviates from the BOM, you may have to use the Async Patch Tool to enable an upgrade to a later version of VMware Cloud Foundation.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy server) for downloading the bundles.

- The computer must have Java 8 or Java 11.
- A Windows or Linux computer with access to the SDDC Manager appliance for uploading the bundles.
- Refer to [KB 88287](#) to ensure that the target version of VMware Cloud Foundation is supported for upgrade. Contact Broadcom Support if you have questions about support for your upgrade path.
- You must have the latest version of the Async Patch Tool.

NOTE

If an existing or older version of the Async Patch Tool exists in the directory, you will need to remove these files from both the Linux or Windows computer and the SDDC manager before downloading the latest version of the Async Patch Tool.

```
rm -r <AP Tool directory>
```

```
rm -r <outputdirectory>
```

The default directory is `/home/vcf/apToolBundles` if `outputDirectory` was not specified when the Async Patch Tool was previously run.

- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Async Patch Tool for long-running operations.
- The Async Patch Tool is supported with VMware Cloud Foundation 4.2.1 and later. This release also supports ESXi and VxRail Manager patching of VMware Cloud Foundation on VxRail.
- You are upgrading from VMware Cloud Foundation 4.x to VMware Cloud Foundation 4.y.

NOTE

This procedure is not required if you are upgrading from VMware Cloud Foundation 4.x to VMware Cloud Foundation 5.0.

After you apply an async patch, and one or more of the VMware Cloud Foundation component versions deviates from the BOM, the process for upgrading to a later version of VMware Cloud Foundation depends on the target version of VMware Cloud Foundation.

Target Version	Upgrade Process
VMware Cloud Foundation 4.y	<p>You must use the Async Patch Tool to enable an upgrade to a later version of VMware Cloud Foundation.</p> <p>For example, if you apply a vCenter Server patch to a VMware Cloud Foundation 4.3.1 instance, you must use the Async Patch Tool to enable upgrade to VMware Cloud Foundation 4.4.</p> <p>The following procedure describes this process.</p>
VMware Cloud Foundation 5.0	<p>Upgrades to 5.0 are automatically enabled and you can upgrade using the Bundle Transfer Utility and SDDC Manager UI.</p> <p>This process is described in the <i>VMware Cloud Foundation Lifecycle Management Guide</i>.</p>

Table continued on next page

Continued from previous page

Target Version	Upgrade Process
	<p>NOTE You should still use the Async Patch Tool to deactivate all async patches and run an inventory sync before upgrading to VMware Cloud Foundation 5.0. See "VCF Async Patch Tool Options" for more information.</p>

If your SDDC Manager appliance does not have a connection to the internet, you can run the Async Patch Tool from a computer that does. Download the upgrade bundles to the computer, copy the bundles to the SDDC Manager appliance, and enable the upgrade using the Async Patch Tool. You can then use the SDDC Manager UI to upgrade VMware Cloud Foundation.

1. Download upgrade bundles.

- a) On the computer with internet access, navigate to `vcf-async-patch-tool-<version>/bin`.
- b) Run the following command:

Linux VMware Cloud Foundation:

```
./vcf-async-patch-tool --download --targetVcfVersion targetVcfVersion --
sourceVcfVersion sourceVcfVersion --sku VCF --depotUser broadcom_support_email
```

Linux VMware Cloud Foundation on Dell EMC VxRail:

```
./vcf-async-patch-tool --download --targetVcfVersion targetVcfVersion --
sourceVcfVersion sourceVcfVersion --sku VCF_ON_VXRAIL --
depotUser broadcom_support_email --pdudell_emc_depot_email
```

Windows VMware Cloud Foundation:

```
vcf-async-patch-tool.bat --download --targetVcfVersion targetVcfVersion --
sourceVcfVersion sourceVcfVersion --sku VCF --depotUser broadcom_support_email
```

Windows VMware Cloud Foundation on Dell EMC VxRail:

```
vcf-async-patch-tool.bat --download --targetVcfVersion targetVcfVersion --
sourceVcfVersion sourceVcfVersion --sku VCF_ON_VXRAIL --depotUser broadcom_support_e
mail --pdudell_emc_depot_email
```

- Replace `targetVcfVersion` with the target version of VMware Cloud Foundation. For example: 4.4.0.0.
- Replace `sourceVcfVersion` with the source version of VMware Cloud Foundation. For example: 4.3.1.0.
- Replace `broadcom_support_email` with your Broadcom Support portal email address.
- Replace `dell_emc_depot_email` with your Dell EMC depot email address. (VxRail only)
- `--outputDirectory` is optional and can be used to specify a location for the download. Choose a directory that has enough free space for the bundles. If you do not specify a location, the Async Patch Tool displays the default location in its output. For example: `/root/apToolBundles`.

NOTE

If you connect to the internet through a proxy server, use the `--proxyServer`, `--ps` option to specify the FQDN and port of the proxy server. For example, `--proxyServer FQDN:port`.

- c) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.
- d) Enter your Broadcom Support portal password.
- e) If the product type is `VX_MANAGER`, enter your Dell EMC Depot user name password. (VxRail only)
- f) Enter `Y` or `N` to choose whether or not to download vRealize bundles.

The Async Patch Tool determines which bundles are required and downloads the bundles .

2. Copy the upgrade bundles and set permissions.

- a) Copy the entire output directory (for example, `apToolBundles`) to the SDDC Manager appliance.

You can choose any location that has enough free space available, for example, `/nfs/vmware/vcf/nfs-mount/apToolBundles`.

- b) SSH in to the SDDC Manager appliance using the `vcf` user account.

- c) Navigate to `/nfs/vmware/vcf/nfs-mount/apToolBundles`.

If you copied the output directory to a different location, navigate to that directory instead.

- d) Run the following commands:

```
chmod -R 755 apToolBundles
chown -R vcf:vcf apToolBundles
```

3. Enable the upgrade.

- a) Navigate to `/home/vcf/asyncPatchTool/bin`.

- b) Run the following command:

```
./vcf-async-patch-tool --enableVCFUpgrade targetVcfVersion --sddcSSOUser ssoUsername --sddcSSHUser vcf --outputDirectory bundleDirectory --it OFFLINE
```

- Replace *targetVcfVersion* with the target version of VMware Cloud Foundation. For example: `4.4.0.0`.
- Replace *ssoUsername* with the management domain SSO user account, for example `administrator@vsphere.local`.
- Replace *bundleDirectory* with the location of the bundle directory from step 2. For example, `/nfs/vmware/vcf/nfs-mount/apToolBundles`.

- c) Enter `Y` to confirm that you are running the latest version of the Async Patch Tool.

- d) Read the information and enter `Y` to acknowledge the pre-requisites.

- e) Enter the password for the super user (`vcf`) account.

- f) Enter the password for the root user account.

- g) Enter the password for the management domain SSO user account.

The Async Patch Tool uploads the upgrade bundles to the internal LCM repository on the SDDC Manager appliance.

4. Log in to the SDDC Manager UI and apply the bundles to all workload domains.

VCF Async Patch Tool Options

The Async Patch Tool is a utility that allows you to apply critical patches outside of the normal VMware Cloud Foundation lifecycle management process. It also provides options for managing async patches and upgrading a VMware Cloud Foundation instance that includes async patches.

Async Patch Tool Help Option

Option	Description
<code>-h, --help</code>	Provides information about the Async Patch Tool options.

Example:

```
./vcf-async-patch-tool -h
```

Customer Experience Improvement Program (CEIP) Option

The Async Patch Tool participates in the Customer Experience Improvement Program (CEIP). You can enable or deactivate CEIP for the Async Patch Tool.

The Customer Experience Improvement Program provides Broadcom with information that enables the company to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, Broadcom collects technical information about your organization's use of the Broadcom products and services regularly in association with your organization's Broadcom license keys. This information does not personally identify any individual. For additional information regarding the CEIP, refer to the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

Option	Description
<code>--ceip true/false</code>	<p>Enable or deactivate telemetry (CEIP) for the Async Patch Tool.</p> <p>NOTE The <code>--ceip</code> option is deprecated. You will be prompted to specify the CEIP settings when running other options using the Async Patch Tool, for example, <code>-l, --listAsyncPatch</code> or <code>-e, --enableAsyncPatch</code>.</p>

Example:

```
./vcf-async-patch-tool -ceip false
```

List Async Patches Option

This option requires an internet connection. Use this option before applying an async patch. See [Apply an Async Patch to VMware Cloud Foundation in Online Mode](#) and [Apply an Async Patch to VMware Cloud Foundation in Offline Mode](#).

Option	Description
<code>-l, --listAsyncPatch</code>	Lists the available async patches. Refer to KB 88287 to see which versions of VMware Cloud Foundation support each async patch.

Required Inputs	Optional Inputs
<p><code>--depotUser, --du</code> Enter your Broadcom Support portal email address for connecting with the VMware Depot.</p>	<ul style="list-style-type: none"> <code>--sku</code> Filters the async patch list by SKU. Enter <code>VCF</code> or <code>VCF_ON_VXRAIL</code>. <code>--productType, --ptype</code> Filters the async patch list by product type. Enter <code>ESX_HOST</code>, <code>NSX</code>, or <code>VCENTER</code>. <p>NOTE Product type <code>ESX_HOST</code> is not available for VMware Cloud Foundation on Dell EMC VxRail.</p> <ul style="list-style-type: none"> <code>--proxyServer, --ps</code>

Table continued on next page

Continued from previous page

Required Inputs	Optional Inputs
	If you connect to the internet through a proxy server, use the <code>--proxyServer</code> , <code>--ps</code> option to specify the FQDN and port of the proxy server. For example, <code>--proxyServer FQDN:port</code> .

Example:

```
./vcf-async-patch-tool --listAsyncPatch --depotUser user@vmware.com --productType VCENTER
```

Download Patch Option (offline only)

This option requires an internet connection. Use this option to download an async patch in an offline environment. See [Apply an Async Patch to VMware Cloud Foundation in Offline Mode](#).

Option	Description
<code>-d</code> , <code>--download</code>	Downloads the specified patch and artifacts.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> <code>--patch</code> Enter the product and version (product:version) of an async patch. For example: VCENTER:7.0.3.00300-19234570. <code>--depotUser</code>, <code>--du</code> Enter your Broadcom Support portal email address for connecting with the VMware Depot. <code>--partnerBundleDepotUserName</code>, <code>--pdu</code> (VxRail only) Enter your Dell EMC depot email address. 	<ul style="list-style-type: none"> <code>--op</code>, <code>--outputDirectory</code> Enter the full path to the location to download the patch. If you do not specify an output directory the Async Patch Tool uses <code>/root/apToolBundles</code>. <code>--sddcManagerVersion</code> Enter your current version of SDDC Manager. For example: 4.5.0.0. This limits the number of bundles that are downloaded to only those that are applicable to your current version of SDDC Manager. <code>--sku</code> Filters the async patch list by SKU. Enter <code>VCF</code> or <code>VCF_ON_VXRAIL</code>. <code>--proxyServer</code>, <code>--ps</code> If you connect to the internet through a proxy server, use the <code>--proxyServer</code>, <code>--ps</code> option to specify the FQDN and port of the proxy server. For example, <code>--proxyServer FQDN:port</code>.

Example:

```
./vcf-async-patch-tool -d --patch VCENTER:7.0.3.00300-19234570 --du user@vmware.com --sku VCF --sddcManagerVersion 4.5.0.0
```

Enable Patch Option

The enable patch option must be run on the SDDC Manager appliance. After you enable an async patch, you can log in to the SDDC Manager UI and apply the patch to all workload domains. See [Apply an Async Patch to VMware Cloud Foundation in Online Mode](#) and [Apply an Async Patch to VMware Cloud Foundation in Offline Mode](#).

Option	Description
<code>-e, --enableAsyncPatch</code>	Enabling a patch performs an inventory sync, bundle download (online mode only), enable patch precheck, and enable patch postcheck. It also uploads the patch to the SDDC Manager appliance internal LCM repository.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> • <code>--patch</code> Enter the product and version (product:version) of an async patch. For example: VCENTER:7.0.3.00300-19234570. • <code>--sddcSSOUser, --ssou</code> Enter the management domain SSO user. For example: administrator@vsphere.local. • <code>--sddcSSHUser, --sshu</code> Enter <code>vcf</code>. • <code>--depotUser, --du</code> (online only) Enter your Broadcom Support portal email address for connecting with the VMware Depot. • <code>--partnerBundleDepotUserName, --pdu</code> (online VxRail only) Enter your Dell EMC depot email address. • <code>--instanceType, --it</code> Enter ONLINE or OFFLINE. 	<ul style="list-style-type: none"> • <code>--op, --outputDirectory</code> Enter the full path to the location to download the patch (online mode) or the location to which you uploaded the patch (offline mode). For example, <code>/nfs/vmware/vcf/nfs-mount/apToolBundles</code>. <p>If you do not specify an output directory the Async Patch Tool uses <code>/root/apToolBundles</code>.</p> <ul style="list-style-type: none"> • <code>--proxyServer, --ps</code> If you connect to the internet through a proxy server, use the <code>--proxyServer, --ps</code> option to specify the FQDN and port of the proxy server. For example, <code>--proxyServer FQDN:port</code>.

Example:

```
./vcf-async-patch-tool -e --patch VCENTER:7.0.3.00300-19234570 --sddcSSOUser
administrator@vsphere.local --sddcSSHUser vcf --outputDirectory /nfs/vmware/vcf/nfs-mount/
apToolBundles --it ONLINE
```

Precheck Option

Prechecks are performed as part of the `-e, --enableAsyncPatch` and `-r, --enableVCFUpgrade` options. You can also run the precheck option on its own prior to enabling an async patch or enabling a VCF upgrade to make sure the operation will succeed. You must run the precheck option on the SDDC Manager appliance.

Option	Description
<code>--pre, --precheck</code>	Validates that the system is able to enable an async patch or enable a VCF upgrade.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> • <code>-e, --enableAsyncPatch</code> or <code>-r, --enableVCFUpgrade</code> For <code>-r, --enableVCFUpgrade</code> you must provide the target version for VCF. For example: <code>4.4.0.0</code>. • <code>--patch</code> (enable patch only) 	<ul style="list-style-type: none"> • <code>--op, --outputDirectory</code> (online only) Enter the full path to the location to download the patch or bundles. For example, <code>/nfs/vmware/vcf/nfs-mount/apToolBundles</code>.

Table continued on next page

Continued from previous page

Required Inputs	Optional Inputs
<p>Enter the product and version (product:version) of an async patch. For example: VCENTER:7.0.3.00300-19234570.</p> <ul style="list-style-type: none"> • <code>--sddcSSOUser</code>, <code>--ssou</code> <p>Enter the management domain SSO user. For example: administrator@vsphere.local.</p> <ul style="list-style-type: none"> • <code>--sddcSSHUser</code>, <code>--sshu</code> <p>Enter <code>vcf</code>.</p> <ul style="list-style-type: none"> • <code>--depotUser</code>, <code>--du</code> (online only) <p>Enter your Broadcom Support portal email address for connecting with the VMware Depot.</p> <ul style="list-style-type: none"> • <code>--partnerBundleDepotUserName</code>, <code>--pdu</code> (online, VxRail, enable VCF upgrade only) <p>Enter your Dell EMC depot email address.</p> <ul style="list-style-type: none"> • <code>--instanceType</code>, <code>--it</code> <p>Enter ONLINE or OFFLINE.</p> <ul style="list-style-type: none"> • <code>--op</code>, <code>--outputDirectory</code> (offline only) <p>Enter the full path to the location to which you uploaded the patch or bundles. For example, <code>/nfs/vmware/vcf/nfs-mount/apToolBundles</code>.</p>	<p>If you do not specify an output directory the Async Patch Tool uses <code>/root/apToolBundles</code>.</p> <ul style="list-style-type: none"> • <code>--proxyServer</code>, <code>--ps</code> <p>If you connect to the internet through a proxy server, use the <code>--proxyServer</code>, <code>--ps</code> option to specify the FQDN and port of the proxy server. For example, <code>--proxyServer FQDN:port</code>.</p>

Online Mode Example:

```
./vcf-async-patch-tool -e --pre --patch VCENTER:7.0.3.00300-19234570 --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf --depotUser user@vmware.com --outputDirectory /nfs/vmware/vcf/nfs-mount/apToolBundles --it ONLINE
```

Offline Mode Example:

```
./vcf-async-patch-tool --pre --enableVCFUpgrade 4.4.0.0 --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf --outputDirectory /nfs/vmware/vcf/nfs-mount/apToolBundles --it OFFLINE
```

Postcheck Option

Postchecks are performed as part of the `-e`, `--enableAsyncPatch` option. You can also run the postcheck option on its own after enabling an async patch. You must run the postcheck option on the SDDC Manager appliance.

Option	Description
<code>--post</code> , <code>--postcheck</code>	Validates that an async patch has been uploaded to the SDDC Manager appliance internal LCM repository and is available for upgrade.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> • <code>-e</code>, <code>--enableAsyncPatch</code> • <code>--patch</code> 	None.

Table continued on next page

Continued from previous page

Required Inputs	Optional Inputs
<p>Enter the product and version (product:version) of an async patch. For example: VCENTER:7.0.3.00300-19234570.</p> <ul style="list-style-type: none"> • <code>--sddcSSOUser</code>, <code>--ssou</code> <p>Enter the management domain SSO user. For example: administrator@vsphere.local.</p> <ul style="list-style-type: none"> • <code>--sddcSSHUser</code>, <code>--sshu</code> <p>Enter <code>vcf</code>.</p> <ul style="list-style-type: none"> • <code>--op</code>, <code>--outputDirectory</code> <p>Enter the full path to the bundle download location that you used when you enabled the patch. For example, <code>/nfs/vmware/vcf/nfs-mount/apToolBundles</code>.</p>	

Example:

```
./vcf-async-patch-tool -e --post --patch VCENTER:7.0.3.00300-19234570 --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf --outputDirectory /nfs/vmware/vcf/nfs-mount/apToolBundles
```

Deactivate All Patches Option

You cannot enable async patches or enable VCF Upgrades if your VMware Cloud Foundation instance already has any async patches enabled. All patches are deactivated when you run the Async Patch Tool with `-r`, `--enableVCFUpgrade` option. You can also run the `--disableAllPatches` option on its own. You must run the option on the SDDC Manager appliance.

Option	Description
<code>--da</code> , <code>--disableAllPatches</code>	Deactivates all async patches on the system that were previously enabled.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> • <code>--sddcSSOUser</code>, <code>--ssou</code> <p>Enter the management domain SSO user. For example: administrator@vsphere.local.</p> <ul style="list-style-type: none"> • <code>--sddcSSHUser</code>, <code>--sshu</code> <p>Enter <code>vcf</code>.</p>	None.

Example:

```
./vcf-async-patch-tool -disableAllPatches --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf
```

Enable VCF Upgrade Option

The enable VCF upgrade option must be run on the SDDC Manager appliance. After you apply an async patch, and one or more of the VMware Cloud Foundation component versions deviates from the BOM, you must use the Async Patch Tool to upgrade from VMware Cloud Foundation 4.x to VMware Cloud Foundation 4.y.

NOTE

If you are upgrading to VMware Cloud Foundation 5.0, you do not need to use the Async Patch Tool to enable upgrade. You should still use the Async Patch Tool to deactivate all async patches and run an inventory sync before upgrading to VMware Cloud Foundation 5.0.

The `-r, --enableVCFUpgrade` option prepares an async patched environment for upgrade and uploads the upgrade bundles to the SDDC Manager appliance internal LCM repository. See [Upgrade an Async Patched Version of VMware Cloud Foundation in Online Mode](#) and [Upgrade an Async Patched Version of VMware Cloud Foundation in Offline Mode](#).

Option	Description
<code>-r, --enableVCFUpgrade</code> Requires the target version for VCF. For example: 4.4.0.0.	Enables upgrade to a target version of VMware Cloud Foundation.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> <code>--sddcSSOUser, --ssou</code> Enter the management domain SSO user. For example: administrator@vsphere.local. <code>--sddcSSHUser, --sshu</code> Enter vcf. <code>--depotUser, --du (online only)</code> Enter your Broadcom Support portal email address for connecting with the VMware Depot. <code>--partnerBundleDepotUserName, --pdu (online VxRail only)</code> Enter your Dell EMC depot email address. <code>--instanceType, --it</code> Enter ONLINE or OFFLINE. <code>--op, --outputDirectory (required for offline only)</code> Enter the full path to the location to which you uploaded the patch or bundles. For example, /nfs/vmware/vcf/nfs-mount/apToolBundles. 	<ul style="list-style-type: none"> <code>--op, --outputDirectory (optional for online)</code> Enter the full path to the location to download the patch or bundles. For example, /nfs/vmware/vcf/nfs-mount/apToolBundles. If you do not specify an output directory the Async Patch Tool uses /root/apToolBundles. <code>--proxyServer, --ps</code> If you connect to the internet through a proxy server, use the <code>--proxyServer, --ps</code> option to specify the FQDN and port of the proxy server. For example, <code>--proxyServer FQDN:port</code>.

Online Mode Example:

```
./vcf-async-patch-tool --enableVCFUpgrade 4.4.0.0 --depotUser user@vmware.com --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf --it ONLINE
```

Offline Mode Example:

```
./vcf-async-patch-tool --enableVCFUpgrade 4.4.0.0 --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf --outputDirectory /nfs/vmware/vcf/nfs-mount/apToolBundles --it OFFLINE
```

Inventory Sync Option

Inventory sync is performed as part of the `-e, --enableAsyncPatch` and `-r, --enableVCFUpgrade` options. You can also run the inventory sync option on its own. You must run the inventory sync option on the SDDC Manager appliance.

This option updates the SDDC Manager inventory with the accurate information about the versions of vCenter Server, NSX, and VMware ESXi that are running in your VMware Cloud Foundation instance. The SDDC Manager inventory can get out of sync if you upgrade any of these components outside of VMware Cloud Foundation.

Option	Description
<code>--performInventorySync, --sync</code>	Updates the SDDC Manager inventory and saves a CSV file with information about each VMware Cloud Foundation component.

Sample output file (in `/home/vcf/asyncPatchTool/bin`):

Software Type	FQDN	Before sync Inventory Version	Current Version	Domain Id(s)	Domain Names	Cluster name
VCENTER	vcenter-1.vrack.vsphere.local	7.0.3.00500-19480866	7.0.3.00500-19480866	[48de0377-98ac-41c8-8dfc-fb916a76af19]	[sddcid-1001]	
NSXT_CLUSTER	vip-nsx-mgmt.vrack.vsphere.local	3.1.3.7.4-19762317	3.1.3.7.4-19762317	[48de0377-98ac-41c8-8dfc-fb916a76af19]	[sddcid-1001]	
ESXI	esxi-1.vrack.vsphere.local	7.0.3-19482537	7.0.3-19482537	[48de0377-98ac-41c8-8dfc-fb916a76af19]	[sddcid-1001]	SDDC-Cluster1
ESXI	esxi-2.vrack.vsphere.local	7.0.3-19482537	7.0.3-19482537	[48de0377-98ac-41c8-8dfc-fb916a76af19]	[sddcid-1001]	SDDC-Cluster1
ESXI	esxi-3.vrack.vsphere.local	7.0.3-19482537	7.0.3-19482537	[48de0377-98ac-41c8-8dfc-fb916a76af19]	[sddcid-1001]	SDDC-Cluster1
ESXI	esxi-4.vrack.vsphere.local	7.0.3-19482537	7.0.3-19482537	[48de0377-98ac-41c8-8dfc-fb916a76af19]	[sddcid-1001]	SDDC-Cluster1

NOTE

In this example, no versions were updated as part of the inventory sync.

Required Inputs	Optional Inputs
<ul style="list-style-type: none"> <code>--sddcSSOUser, --ssou</code> Enter the management domain SSO user. For example: <code>administrator@vsphere.local</code>. <code>--sddcSSHUser, --sshu</code> Enter <code>vcf</code>. 	None.

Example:

```
./vcf-async-patch-tool --sync --sddcSSOUser administrator@vsphere.local --sddcSSHUser vcf
```

VCF on Dell VxRail Guide

Provides information about deploying, managing, and upgrading VMware Cloud Foundation on Dell VxRail.

About VMware Cloud Foundation on Dell VxRail

The *VMware Cloud Foundation on Dell VxRail Guide* provides information on managing the integration of VMware Cloud Foundation and Dell VxRail. As this product is an integration of VMware Cloud Foundation and Dell VxRail, the expected results are obtained only when the configuration is done from both the products. This guide covers all the information regarding the VMware Cloud Foundation workflows. For the instructions on configuration to be done on Dell VxRail, this guide provides links to the Dell VxRail documentation.

Intended Audience

The *VMware Cloud Foundation on Dell VxRail Guide* is intended for the system administrators of the VxRail environments who want to adopt VMware Cloud Foundation. The information in this document is written for experienced data center system administrators who are familiar with:

- Concepts of virtualization, software-defined data centers, and virtual infrastructure (VI).
- VMware virtualization technologies, such as VMware ESXi™, the hypervisor
- Software-defined networking using VMware NSX®
- Software-defined storage using VMware vSAN™
- IP networks

Additionally, you should be familiar with these software products, software components, and their features:

- Dell EMC VxRail Manager
- VMware vSphere®
- VMware vCenter Server® and VMware vCenter Server®Appliance™
- VMware vRealize®Log Insight™
- VMware vSphere® with VMware Tanzu™

Related Publications

The *Planning and Preparation Workbook* provides detailed information about the software, tools, and external services that are required to deploy VMware Cloud Foundation on Dell EMC VxRail.

The *VMware Cloud Foundation on Dell Release Notes* provide information about each release, including:

- What's new in the release
- Software components and versions included in the Bill of Materials (BOM)
- Resolved issues
- Known issues

The *VMware Cloud Foundation on Dell VxRail API Reference Guide* provides information about using the API.

VMware Cloud Foundation on Dell VxRail

VMware Cloud Foundation on Dell VxRail enables VMware Cloud Foundation on top of the Dell VxRail platform.

An administrator of a VMware Cloud Foundation on Dell VxRail system performs tasks such as:

- Deploy VMware Cloud Foundation on Dell VxRail.
- Manage certificates.
- Add capacity to your system.
- Configure and provision workload domains.

- Manage provisioned workload domains.
- Monitor alerts and the health of the system.
- Troubleshoot issues and prevent problems across the physical and virtual infrastructure.
- Perform life cycle management on the software components.

Prepare a VxRail Environment for Cloud Builder Appliance Deployment

Before you can deploy the VMware Cloud Builder Appliance on the VxRail cluster, you must complete the following tasks.

Imaging the VxRail Management Nodes

Image the VxRail management nodes by using Dell RASR (Rapid Appliance Self Recovery) process. Ensure that you update the RASR image in each server node SD card before you start the imaging process.

For detailed information about how to image the VxRail management nodes, contact Dell Support.

VxRail First Run for the Management Cluster

Before you can deploy the management domain using VMware Cloud Builder, you must perform the VxRail first run.

The VxRail first run for the management cluster consists of the following tasks:

- The discovery of the VxRail Nodes occurs. All the nodes that were imaged are detected.
- Upload the JSON configuration file. Trigger the validation.
- All the configuration inputs are validated.

NOTE

You specify the vSphere Lifecycle Manager method (vSphere Lifecycle Manager images or vSphere Lifecycle Manager baselines) during the VxRail first run or in the vCenter Server VxRail UI after the first run, but before bring-up. See the Dell VxRail documentation for details.

NOTE

vSAN Express Storage Architecture (ESA) requires vSphere Lifecycle Manager images.

The following components are deployed and enabled:

- vCenter
- vSAN
- VxRail Manager

Click **Manage VxRail** to log in to the VMware vCenter server.

For information on VxRail First Run, contact Dell Support.

Deploy VMware Cloud Builder Appliance

VMware Cloud Builder is a virtual appliance that is used to deploy and configure the first cluster of the management domain and transfer inventory and control to SDDC Manager. During the deployment process, the VMware Cloud Builder appliance validates network information you provide in the deployment parameter workbook such as DNS, network (VLANS, IPs, MTUs), and credentials.

Before you deploy the VMware Cloud Builder appliance, verify that your environment fulfills the requirements for this process.

Prerequisite	Value
Environment	<ul style="list-style-type: none"> • Verify that your environment is configured for deployment of VMware Cloud Builder and the management domain. • Verify that you have available virtual infrastructure that has access to the management network that will be

Table continued on next page

Continued from previous page

Prerequisite	Value
	used by the management domain. You deploy VMware Cloud Builder on that virtual infrastructure.
Resource Requirements	<ul style="list-style-type: none"> • 4 CPUs • 4 GB of Memory • 279 GB of Storage <ul style="list-style-type: none"> – 25.1 GB (thin provisioned) – 253.8 GB (thick provisioned)
Installation Packages	Verify that you download the OVA file(s) for VMware Cloud Builder.
Network	<ul style="list-style-type: none"> • Verify that the static IP address and FQDN for the VMware Cloud Builder appliance are available. • Verify that connectivity is in place from the VMware Cloud Builder appliance and the management VLAN used in the deployment.

The VMware Cloud Builder appliance must be on the same management network as the hosts to be used. It must also be able to access all required external services, such as DNS and NTP.

This procedure describes deploying the VMware Cloud Builder appliance to the cluster that was created during the VxRail first run.

1. Download the VMware Cloud Builder appliance OVA.
2. Log in to vCenter Server using the vSphere Client.
3. In the navigator, select the cluster that was created during the VxRail first run.
4. Click **Actions** > **Deploy OVF Template**.
5. Select **Local file** and click **Upload Files**.
6. Browse to the VMware Cloud Builder appliance OVA, select it, and click **Open**.
7. Click **Next**.
8. Enter a name for the virtual machine, select a target location, and click **Next**.
9. Select the cluster you created during the VxRail first run and click **Next**.
10. Review the details and click **Next**.
11. Accept the license agreement and click **Next**.
12. On the Select Storage page, select the storage for the VMware Cloud Builder appliance and click **Next**.
13. On the Select networks dialog box, select the management network and click **Next**.
14. On the Customize template page, enter the following information for the VMware Cloud Builder appliance and click **Next**:

Setting	Details
Admin Username	The admin user name cannot be one of the following pre-defined user names: <ul style="list-style-type: none"> • root • bin • daemon • messagebus • systemd-bus-proxy • systemd-journal-gateway • systemd-journal-remote

Table continued on next page

Continued from previous page

Setting	Details
	<ul style="list-style-type: none"> • systemd-journal-upload • systemd-network • systemd-resolve • systemd-timesync • nobody • sshd • named • rpc • tftp • ntp • smmsp • cassandra
Admin Password/Admin Password confirm	<p>The admin password must be a minimum of 15 characters and include at least one uppercase, one lowercase, one digit, and one special character. Supported special characters:</p> <p>@ ! # \$ % ? ^</p> <p>NOTE A password cannot be based on a dictionary word (for example, VMware1!)</p>
Root password/Root password confirm	<p>The root password must be a minimum of 15 characters and include at least one uppercase, one lowercase, one digit, and one special character. Supported special characters:</p> <p>@ ! # \$ % ? ^</p> <p>NOTE A password cannot be based on a dictionary word (for example, VMware1!)</p>
Hostname	Enter the hostname for the VMware Cloud Builder appliance.
Network 1 IP Address	Enter the IP address for the VMware Cloud Builder appliance.
Network 1 Subnet Mask	For example, 255 . 255 . 255 . 0.
Default Gateway	Enter the default gateway for the VMware Cloud Builder appliance.
DNS Servers	IP address of the primary and secondary DNS servers (comma separated). Do not specify more than two servers.
DNS Domain Name	For example, vsphere . local.
DNS Domain Search Paths	Comma separated. For example vsphere . local , sf . vsphere . local.
NTP Servers	Comma separated.

15. Review the deployment details and click **Finish**.

NOTE

Make sure your passwords meet the requirements specified above before clicking **Finish** or your deployment will not succeed.

16. After the VMware Cloud Builder appliance is deployed, SSH in to the VM with the admin credentials provided in step 14.
17. Ensure that you can ping the ESXi hosts.
18. Verify that the VMware Cloud Builder appliance has access to the required external services, such as DNS and NTP by performing forward and reverse DNS lookups for each host and the specified NTP servers.

Deploy the Management Domain Using VMware Cloud Builder

The VMware Cloud Foundation deployment process is referred to as bring-up. You specify deployment information specific to your environment such as networks, hosts, license keys, and other information in the deployment parameter workbook and upload the file to the VMware Cloud Builder appliance to initiate bring-up of the management domain.

During bring-up, the management domain is created on the ESXi hosts specified in the deployment parameter workbook. The VMware Cloud Foundation software components are automatically deployed, configured, and licensed using the information provided. The deployment parameter workbook can be reused to deploy multiple VMware Cloud Foundation instances of the same version.

The following procedure describes how to perform bring-up of the management domain using the deployment parameter workbook. You can also perform bring-up using a custom JSON specification. See the [VMware Cloud Foundation API Reference Guide](#) for more information.

Externalizing the vCenter Server that gets created during the VxRail first run is automated as part of the bring-up process.

1. In a web browser, log in to the VMware Cloud Builder appliance administration interface: `https://Cloud_Builder_VM_FQDN`.
2. Enter the admin credentials you provided when you deployed the VMware Cloud Builder appliance and then click **Log In**.
3. On the **End-User License Agreement** page, select the **I Agree to the End User License Agreement** check box and click **Next**.
4. Select **VMware Cloud Foundation on Dell EMC VxRail** and click **Next**.
5. Review and acknowledge the prerequisites and click **Next**.

If there are any gaps, ensure they are fixed before proceeding to avoid issues during the bring-up process. You can download or print the prerequisite list for reference.

6. Download the deployment parameter workbook from the [Broadcom Support portal](#) and fill it in with the required information.
See [About the Deployment Parameter Workbook](#).

7. Click **Next**.
8. Click **Select File**, browse to the completed workbook, and click **Open** to upload the workbook.
9. Click **Next** to begin validation of the uploaded file.

To access the bring-up log file, SSH to the VMware Cloud Builder appliance as `admin` and open the `/opt/vmware/bringup/logs/vcf-bringup-debug.log` file.

If there is an error during the validation and the **Next** button is grayed out, you can either make corrections to the environment or edit the deployment parameter workbook and upload it again. Then click **Retry** to perform the validation again.

If any warnings are displayed and you want to proceed, click **Acknowledge** and then click **Next**.

10. Click **Deploy SDDC**.

During the bring-up process, the vCenter Server, NSX, and SDDC Manager appliances are deployed and the management domain is created. The status of the bring-up tasks is displayed in the UI.

After bring-up is completed, a green bar is displayed indicating that bring-up was successful. A link to the SDDC Manager UI is also displayed. If there are errors during bring-up, see [Troubleshooting Deployment](#).

11. Click **Download** to download a detailed deployment report. This report includes information on assigned IP addresses and networks that were configured in your environment.
12. After bring-up is completed, click **Finish**.
13. Click **Launch SDDC Manager**.
14. Power off the VMware Cloud Builder appliance.

About the Deployment Parameter Workbook

The deployment parameter workbook contains worksheets categorizing the information required for deploying VMware Cloud Foundation. The information provided is used to create the management domain using the VMware Cloud Builder appliance.

Before you begin filling in the deployment parameter workbook, download the workbook from the [Broadcom Support portal](#).

The fields in yellow contain sample values that you should replace with the information for your environment. If a cell turns red, the required information is missing, or validation input has failed.

IMPORTANT

The deployment parameter workbook is not able to fully validate all inputs due to formula limitations of Microsoft Excel. Some validation issues may not be reported until you upload the deployment parameter workbook to the VMware Cloud Builder appliance.

NOTE

Do not copy and paste content between cells in the deployment parameter workbook, since this may cause issues.

The Introduction worksheet in the deployment parameter workbook contains an overview of the workbook and guidance on how to complete it. For information about the prerequisites for deploying the management domain, see the *Planning and Preparation Workbook*.

VxRail Prerequisites

- The VxRail first run is completed and vCenter Server and VxRail Manager VMs are deployed.
- The vCenter Server version matches the build listed in the Cloud Foundation Bill of Materials (BOM). See the *VMware Cloud Foundation Release Notes* for the BOM.

Credentials Worksheet

The Credentials worksheet details the accounts and initial passwords for the VMware Cloud Foundation components. You must provide input for each yellow box. A red cell may indicate that validations on the password length has failed.

Input Required

Update the Default Password field for each user (including the automation user in the last row). Passwords can be different per user or common across multiple users. The tables below provide details on password requirements.

Table 216: Password Complexity

Password	Requirements
VxRail Manager root account	Standard
VxRail Manager service account (mystic)	Standard. The service account password must be different than the VxRail Manager root account password.
ESXi Host root account	This is the password which you configured on the hosts during ESXi installation.
Default Single-Sign on domain administrator user	<ol style="list-style-type: none"> Length 8-20 characters Must include: <ul style="list-style-type: none"> – mix of upper-case and lower-case letters – a number – a special character, such as @ ! # \$ % ^ or ? Must not include * { } [] () / \ ' " ` ~ , ; : . < >
vCenter Server virtual appliance root account	<ol style="list-style-type: none"> Length 8-20 characters Must include: <ul style="list-style-type: none"> – mix of upper-case and lower-case letters – a number – a special character, such as @ ! # \$ % ^ or ? Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
NSX virtual appliance root account	<ol style="list-style-type: none"> Length 12-127 characters Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? – at least five different characters Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
NSX user interface and default CLI admin account	<ol style="list-style-type: none"> Length 12-127 characters Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? – at least five different characters Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
NSX audit CLI account	<ol style="list-style-type: none"> Length 12-127 characters Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? – at least five different characters Must not include: * { } [] () / \ ' " ` ~ , ; : . < >
SDDC Manager appliance root account	<ol style="list-style-type: none"> Minimum length 15 characters Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ?

Table continued on next page

Continued from previous page

Password	Requirements
	3. Must not include: <ul style="list-style-type: none"> – *{}[]()/'"``~ , ; : . < > – A dictionary word (for example, VMware1!)
SDDC Manager super user (vcf)	1. Minimum length 15 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? 3. Must not include: <ul style="list-style-type: none"> – *{}[]()/'"``~ , ; : . < > – A dictionary word (for example, VMware1!)
SDDC Manager local account (admin@local)	1. Length 12-127 characters 2. Must include: <ul style="list-style-type: none"> – mix of uppercase and lowercase letters – a number – a special character, such as @ ! # \$ % ^ or ? 3. Must not include: *{}[]()/'"``~ , ; : . < >

Hosts and Networks Worksheet

The Hosts and Networks worksheet specifies the details for all networks and hosts. This information is configured on the appropriate VMware Cloud Foundation components.

Management Domain Networks

This section covers the VLANs, gateways, MTU, and expected IP ranges and subnet mask for each network you have configured on the Top of Rack switches in your environment.

With VMware Cloud Foundation 5.1 and later, you have the ability to create separate distributed port groups for management VM (for example, vCenter Server and NSX Manager) traffic and ESXi host management traffic. You can configure this during the VxRail first run.

Network Type	VLAN	Portgroup Name	CIDR Notation	Gateway	MTU
Management Network	Enter the VLAN ID. The VLAN ID can be between 0 and 4094.	You cannot change the portgroup name prefix.	Enter the CIDR notation for the management network only.	Enter the gateway IP for the management network only.	Enter MTU for the management network only.

Table continued on next page

Continued from previous page

Network Type	VLAN	Portgroup Name	CIDR Notation	Gateway	MTU
vMotion Network	NOTE The VLAN ID for Uplink 1 and Uplink 2 Networks must be unique and not used by any other network type.		NOTE VxRail Manager configures the vMotion and vSAN networks.	NOTE VxRail Manager configures the vMotion and vSAN networks.	NOTE VxRail Manager configures the vMotion and vSAN networks. The MTU can be between 1500 and 9000.
vSAN Network					

System vSphere Distributed Switch Used for NSX Overlay and VLAN Traffic

In VxRail Manager, you can choose to create one or two vSphere Distributed Switches (vDS) for system traffic and to map physical NICs (pNICs) to those vSphere Distributed Switches. The following fields are used to specify which system vDS and vmnics to use for NSX traffic (NSX Overlay, NSX VLAN, Edge Overlay, and Uplink networks). You can also choose to create two additional vDSes to use for NSX traffic. The Transport Zone Type indicates the type of NSX traffic the vDS will be associated with (Overlay, VLAN, or Overlay/VLAN).

NOTE

At least one vDS needs to be marked for Overlay.

System vSphere Distributed Switch - Name	Enter the name of the vDS to use for overlay traffic.
System vSphere Distributed Switch - vmnics to be used for overlay traffic	Enter the vmnics to use for overlay traffic.
System vSphere Distributed Switch - Transport Zone Type	Select Overlay, VLAN, or Overlay/VLAN.

Secondary System vSphere Distributed Switch for NSX Overlay and VLAN Traffic

Choose **Yes** to use a secondary system vDS for overlay/VLAN traffic.

Secondary System vSphere Distributed Switch - Name	Enter the name of the secondary system vSphere Distributed Switch (vDS).
Secondary System vSphere Distributed Switch - vmnics	Enter the vmnics to assign to the secondary system vDS. For example: vmnic4, vmnic5
Secondary System vSphere Distributed Switch - Transport Zone Type	Select Overlay, VLAN, or Overlay/VLAN.

Create Separate vSphere Distributed Switch for NSX Overlay/VLAN Traffic

If you want to use one of the system vSphere Distributed Switches that you created in VxRail Manager for overlay traffic (Host Overlay, Edge Overlay, and Uplink networks), choose **No**. Choose **Yes** to create a new vDS for overlay/VLAN traffic.

New vSphere Distributed Switch - Name	Enter a name for the new vSphere Distributed Switch (vDS).
New vSphere Distributed Switch - vmnics	Enter the vmnics to assign to the new vDS. For example: vmnic4, vmnic5
New vSphere Distributed Switch - MTU Size	Enter the MTU size for the new vDS. Default value is 9000.
New vSphere Distributed Switch - Transport Zone Type	Select Overlay, VLAN, or Overlay/VLAN.

Management Domain ESXi Hosts

Specify the IP addresses of the ESXi hosts for the management domain. In a standard deployment, only four hosts are required in the management domain. VMware Cloud Foundation can also be deployed with a consolidated architecture. In a consolidated deployment, all workloads are deployed in the management domain instead of to separate workload domains. As such, additional hosts may be required to provide the capacity needed. In this section, only enter values for the number of hosts desired in the management domain.

Host Name	IP Address
Enter host names for each of the four ESXi hosts.	Enter IP Address for each of the four ESXi hosts.

ESXi Host Security Thumbprints

If you want bring-up to validate the SSH fingerprints of the ESXi hosts and the SSH fingerprint and SSL thumbprint of the vCenter Server and VxRail Manager to reduce the chance of Man In The Middle (MiTM) attack, select **Yes** in the **Validate Thumbprints** field.

If you set **Validate Thumbprints** to **Yes**, follow the steps below.

1. In a web browser, log in to the ESXi host using the VMware Host Client.
2. In the navigation pane, click **Manage** and click the **Services** tab.
3. Select the **TSM-SSH** service and click **Start** if not started.
4. Connect to the VMware Cloud Builder appliance using an SSH client such as Putty.
5. Enter the admin credentials you provided when you deployed the VMware Cloud Builder appliance.
6. Retrieve the ESXi SSH fingerprints by entering the following command replacing *hostname* with the FQDN of the first ESXi host:

```
ssh-keygen -lf <(ssh-keyscan hostname 2>/dev/null
```
7. In the VMware Host Client, select the **TSM-SSH** service for the ESXi host and click **Stop**.
8. Repeat for the remaining ESXi hosts.
9. Retrieve the vCenter Server SSH fingerprint by entering the following command replacing *hostname* with the FQDN of your vCenter Server:

```
ssh-keygen -lf <(ssh-keyscan hostname 2>/dev/null
```
10. Retrieve the vCenter Server SSL thumbprint by entering the following command replacing *hostname* with the FQDN of your vCenter Server:

```
openssl s_client -connect hostname:443 < /dev/null 2> /dev/null | openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```
11. Retrieve the VxRail Manager SSH fingerprint by entering the following command replacing *hostname* with the FQDN of your VxRail Manager:

```
ssh-keygen -lf <(ssh-keyscan hostname 2>/dev/null
```
12. Retrieve the VxRail Manager SSL thumbprint by entering the following command replacing *hostname* with the FQDN of your VxRail Manager:

```
openssl s_client -connect hostname:443 < /dev/null 2> /dev/null | openssl x509 -sha256 -fingerprint -noout -in /dev/stdin
```

13. Enter the information in the deployment parameter workbook.

NSX Host Overlay Network

By default, VMware Cloud Foundation uses DHCP for the management domain Host Overlay Network TEPs. For this option, a DHCP server must be configured on the NSX host overlay (Host TEP) VLAN of the management domain. When NSX creates TEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.

For the management domain and VI workload domains with uniform L2 clusters, you can choose to use static IP addresses instead. Make sure the IP range includes enough IP addresses for the number of hosts that will use the static IP Pool. The number of IP addresses required depends on the number of pNICs on the ESXi hosts that are used for the vSphere Distributed Switch that handles host overlay networking. For example, a host with four pNICs that uses two pNICs for host overlay traffic requires two IP addresses in the static IP pool..

Table 217: DHCP Settings

Parameter	Value
VLAN ID	Enter a VLAN ID for the NSX host overlay network. The VLAN ID can be between 0 and 4094.
Configure NSX Host Overlay Using a Static IP Pool	Select No to use DHCP.

Table 218: Static IP Pool Settings

Parameter	Value
VLAN ID	Enter a VLAN ID for the NSX host overlay network. The VLAN ID can be between 0 and 4094.
Configure NSX Host Overlay Using a Static IP Pool	Select Yes to use a static IP pool.
Pool Description	Enter a description for the static IP pool.
Pool Name	Enter a name for the static IP pool.
CIDR Notation	Enter CIDR notation for the NSX Host Overlay network.
Gateway	Enter the gateway IP address for the NSX Host Overlay network.
NSX Host Overlay Start IP	Enter the first IP address to include in the static IP pool.
NSX Host Overlay End IP	Enter the last IP address to include in the static IP pool.

Deploy Parameters Worksheet: Existing Infrastructure Details

Your existing DNS infrastructure is used to provide forward and reverse name resolution for all hosts and VMs in the VMware Cloud Foundation SDDC. External NTP sources are also utilized to synchronize the time between the software components.

Table 219: Infrastructure

Parameter	Value
DNS Server #1	Enter IP address of first DNS server.
DNS Server #2	Enter IP address of second DNS server. NOTE If you have only one DNS server, enter n/a in this cell.

Table continued on next page

Continued from previous page

Parameter	Value
NTP Server #1	Enter IP address or FQDN of first NTP server.
NTP Server #2	Enter IP address or FQDN of second NTP server. NOTE If you have only one NTP server, enter n/a in this cell.

Table 220: DNS Zone

Parameter	Value
DNS Zone Name	Enter root domain name for your SDDC management components. NOTE VMware Cloud Foundation expects all components to be part of the same DNS zone.

Table 221: Customer Experience Improvement Program

Parameter	Value
Enable Customer Experience Improvement Program ("CEIP")	Select an option to activate or deactivate CEIP across vSphere, NSX, and vSAN during bring-up.

Table 222: Enable FIPS Security Mode on SDDC Manager

Parameter	Value
Enable FIPS Security Mode on SDDC Manager	Select an option to activate or deactivate FIPS security mode during bring-up. VMware Cloud Foundation supports Federal Information Processing Standard (FIPS) 140-2. FIPS 140-2 is a U.S. and Canadian government standard that specifies security requirements for cryptographic modules. When you enable FIPS compliance, VMware Cloud Foundation enables FIPS cipher suites and components are deployed with FIPS enabled. To learn more about support for FIPS 140-2 in VMware products, see https://www.vmware.com/solutions/security/certifications/fips . NOTE This option is only available for new VMware Cloud Foundation installations and the setting you apply during bring-up will be used for future upgrades. You cannot change the FIPS security mode setting after bring-up.

Deploy Parameters Worksheet: VxRail Manager Details

The VxRail Manager Details section of the Deploy Parameters Worksheet specifies the details for VxRail Manager.

VxRail Manager Details

Enter a host name and an IP address for VxRail Manager.

Deployment Parameters Worksheet: License Keys

Provide licensing information for VMware Cloud Foundation.

1. Select **Yes** or **No** for **License Now**.
2. If you select **Yes**, in the License Keys section, update the red fields with your license keys. Ensure the license key matches the product listed in each row and that the license key is valid for the version of the product listed in the VMware Cloud Foundation BOM. The license key audit during bring-up validates both the format and validity of the key.

NOTE

When using the per-TiB license for vSAN, be aware that VI workload domain components like vCenter and NSX Manager will also consume the TiB capacity.

3. If you select **No**, the VMware Cloud Foundation components are deployed in evaluation mode.

IMPORTANT

After bring-up, you must switch to licensed mode by adding component license keys in the SDDC Manager UI or adding and assigning a solution license key in the vSphere Client. See the *VMware Cloud Foundation Administration Guide* for information about adding component license keys in the SDDC Manager UI. See [Managing vSphere Licenses](#) for more information about adding and applying a solution license key for VMware ESXi and vCenter Server in the vSphere Client. If you are using a solution license key, you must also add a separate VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

Deploy Parameters Worksheet: vSphere Infrastructure

The vSphere infrastructure section of the Deploy Parameters Worksheet details how you want to configure the vCenter Server and its related objects.

This section of the deployment parameter workbook contains sample configuration information, but you can update them with names that meet your naming standards.

NOTE

All host names entries within the deployment parameter workbook expect the short name. VMware Cloud Builder takes the host name and the DNS zone provided to calculate the FQDN value and performs validation prior to starting the deployment. The specified host names and IP addresses must be resolvable using the DNS servers provided, both forward (hostname to IP) and reverse (IP to hostname), otherwise the bring-up process will fail.

Table 223: vCenter Server

Parameter	Host Name	IP Address
vCenter Server	Enter a host name for the vCenter Server.	Enter the IP address for the vCenter Server that is part of the management VLAN. NOTE This is the same VLAN and IP address space where the ESXi management VMKernels reside.

Table 224: vCenter Datacenter and Cluster

Parameter	Value
Datacenter Name	Enter a name for the management datacenter.
Cluster Name	Enter a name for the management cluster.

NOTE

You specify the vSphere Lifecycle Manager method (vSphere Lifecycle Manager images or vSphere Lifecycle Manager baselines) for the vSAN cluster during the VxRail first run. vSAN Express Storage Architecture (ESA) requires vSphere Lifecycle Manager images.

NOTE

Enhanced vMotion Compatibility (EVC) is automatically enabled on the VxRail management cluster.

Select the architecture model you plan to use. If you choose **Consolidated**, specify the names for the vSphere resource pools. You do not need to specify resource pool names if you are using the standard architecture model. See *Introducing VMware Cloud Foundation* for more information about these architecture models.

Table 225: vSphere Resource Pools

Parameter	Value
Resource Pool SDDC Management	Specify the vSphere resource pool name for management VMs.
Resource Pool User Edge	Specify the vSphere resource pool name for user deployed NSX VMs in a consolidated architecture.
Resource Pool User VM	Specify the vSphere resource pool name for user deployed workload VMs.

NOTE

Resource pools are created with Normal CPU and memory shares.

Table 226: vSphere Datastore

Parameter	Value
vSAN Datastore Name	Enter vSAN datastore name for your management components.

NOTE

You specify the vSAN storage architecture (vSAN ESA or vSAN OSA) during the VxRail first run. To use vSAN Express Storage Architecture (ESA) you must use vSphere Lifecycle Manager images for managing the lifecycle of ESXi hosts in the primary cluster of management domain.

If the VMware Cloud Builder appliance does not have direct internet access, you can configure a proxy server to download the vSAN HCL JSON. A recent version of the HCL JSON file is required for vSAN ESA.

Table 227: Proxy Server Configuration

Parameter	Value
Proxy Server Configuration	Select Yes to configure a proxy server.

Table continued on next page

Continued from previous page

Parameter	Value
Proxy Server	Enter the proxy server FQDN or IP address.
Proxy Port	Enter the proxy server port.
Proxy Username	
Proxy Password	
Proxy Transfer Protocol	
HTTPs Proxy Certificate (PEM Encoded)	

Deploy Parameters Worksheet: VMware NSX

The NSX section of the Deploy Parameters Worksheet specifies the details you want to use for deploying VMware NSX components.

Table 228: NSX Management Cluster

Parameter	Value
NSX Management Cluster VIP	Enter the host name and IP address for the NSX Manager VIP. The host name can match your naming standards but must be registered in DNS with both forward and reverse resolution matching the specified IP. NOTE This is the same VLAN and IP address space where the vCenter and ESXi management VMKernels reside.
NSX Virtual Appliance Node #1	Enter the host name and IP address for the first node in the NSX Manager cluster.
NSX Virtual Appliance Node #2	Enter the host name and IP address for the second node in the NSX Manager cluster.
NSX Virtual Appliance Node #3	Enter the host name and IP address for the third node in the NSX Manager cluster.
NSX Virtual Appliance Size	Select the size for the NSX Manager virtual appliances. The default is medium.

Deploy Parameters Worksheet: SDDC Manager

The SDDC Manager section of the Deploy Parameters Worksheet specifies the details for deploying SDDC Manager.

Table 229: SDDC Manager

Parameter	Value
SDDC Manager Hostname	Enter a host name for the SDDC Manager VM.
SDDC Manager IP Address	Enter an IP address for the SDDC Manager VM.

Table continued on next page

Continued from previous page

Parameter	Value
Cloud Foundation Management Domain Name	Enter a name for the management domain. This name will appear in Inventory > Workload Domains in the SDDC Manager UI.

Troubleshooting VMware Cloud Foundation Deployment

During the deployment stage of VMware Cloud Foundation you can use log files and the Supportability and Serviceability (SoS) Tool to help with troubleshooting.

Using the SoS Utility on VMware Cloud Builder

You can run the Supportability and Serviceability (SoS) Utility on the VMware Cloud Builder appliance to generate a support bundle, which you can use to help debug a failed bring-up of VMware Cloud Foundation.

NOTE

After a successful bring-up, you should only run the SoS Utility on the SDDC Manager appliance. See [Supportability and Serviceability \(SoS\) Utility](#) in the *VMware Cloud Foundation Administration Guide*.

The SoS Utility is not a debug tool, but it does provide health check operations that can facilitate debugging a failed deployment.

To run the SoS Utility in VMware Cloud Builder, SSH in to the VMware Cloud Builder appliance using the `admin` administrative account, then enter `su` to switch to the root user, and navigate to the `/opt/vmware/sddc-support` directory and type `./sos` followed by the options required for your desired operation.

```
./sos --option-1--option-2... --option-n
```

SoS Utility Help Options

Use these options to see information about the SoS tool itself.

Option	Description
<code>--help</code> <code>-h</code>	Provides a summary of the available SoS tool options
<code>--version</code> <code>-v</code>	Provides the SoS tool's version number.

SoS Utility Generic Options

These are generic options for the SoS Utility.

Option	Description
<code>--configure-sftp</code>	Configures SFTP for logs.
<code>--debug-mode</code>	Runs the SoS tool in debug mode.
<code>--force</code>	Allows SoS operations from the VMware Cloud Builder appliance after bring-up.

Table continued on next page

Continued from previous page

Option	Description
	<p>NOTE</p> <p>In most cases, you should not use this option. Once bring-up is complete, you can run the SoS Utility directly from the SDDC Manager appliance.</p>
--history	Displays the last twenty SoS operations performed.
--log-dir <i>LOGDIR</i>	Specifies the directory to store the logs.
--log-folder <i>LOGFOLDER</i>	Specifies the name of the log directory.
--setup-json <i>SETUP_JSON</i>	<p>Custom setup-json file for log collection.</p> <p>SoS prepares the inventory automatically based on the environment where it is running. If you want to collect logs for a pre-defined set of components, you can create a <code>setup.json</code> file and pass the file as input to SoS. A sample JSON file is available on the VMware Cloud Builder in the <code>/opt/vmware/sddc-support/</code> directory.</p>
--skip-known-host-check	Skips the specified check for SSL thumbprint for host in the known host.
--zip	Creates a zipped tar file for the output.

SoS Utility Log File Options

Option	Description
--api-logs	Collects output from APIs.
--cloud-builder-logs	Collects Cloud Builder logs.
--esx-logs	<p>Collects logs from the ESXi hosts only.</p> <p>Logs are collected from each ESXi host available in the deployment.</p>
--no-clean-old-logs	<p>Use this option to prevent the tool from removing any output from a previous collection run.</p> <p>By default, before writing the output to the directory, the tool deletes the prior run's output files that might be present. If you want to retain the older output files, specify this option.</p>
--no-health-check	Skips the health check executed as part of log collection.
--nsx-logs	Collects logs from the NSX Manager instances only.
--rvc-logs	<p>Collects logs from the Ruby vSphere Console (RVC) only. RVC is an interface for ESXi and vCenter.</p> <p>NOTE</p> <p>If the Bash shell is not enabled in vCenter, RVC log collection will be skipped .</p> <p>NOTE</p> <p>RVC logs are not collected by default with <code>./sos log collection</code>.</p>
--sddc-manager-logs	Collects logs from the SDDC Manager only.
--test	Collects test logs by verifying the files.

Table continued on next page

Continued from previous page

Option	Description
--vc-logs	Collects logs from the vCenter Server instances only. Logs are collected from each vCenter server available in the deployment.
--vm-screenshots	Collects screen shots from all VMs.

SoS Utility JSON Generator Options

The JSON generator options within the SoS Utility provide a method to execute the creation of the JSON file from a completed deployment parameter workbook. To run the JSON generator, you must provide, as a minimum, a path to the deployment parameter workbook and the design type using the following syntax:

```
./sos --jsongenerator --jsongenerator-input JSONGENERATORINPUT --jsongenerator-design JSONGENERATORDESIGN
```

Option	Description
--jsongenerator	Invokes the JSON generator utility.
--jsongenerator-input <i>JSONGENERATORINPUT</i>	Specify the path to the input file to be used by the JSON generator utility. For example: /tmp/vcf-ems-deployment-parameter.xlsx.
--jsongenerator-design <i>JSONGENERATORDESIGN</i>	Use vcf-vxrail for VMware Cloud Foundation on Dell VxRail.
--jsongenerator-supress	Supress confirmation to force cleanup directory. (optional)
--jsongenerator-logs <i>JSONGENERATORLOGS</i>	Set the directory to be used for logs. (optional)

SoS Utility Health Check Options

The SoS Utility can be used to perform health checks on various components or services, including connectivity, compute, and storage.

NOTE

The health check options are primarily designed to run on the SDDC Manager appliance. Running them on the VMware Cloud Builder appliance requires the `--force` parameter, which instructs the SoS Utility to identify the SDDC Manager appliance deployed by VMware Cloud Builder during the bring-up process, and then execute the health check remotely. For example:

```
./sos --health-check --force
```

Option	Description
--certificate-health	Verifies that the component certificates are valid (within the expiry date).
--connectivity-health	Performs a connectivity health check to inspect whether the different components of the system such as the ESXi hosts, vCenter Servers, NSX Manager VMs, and SDDC Manager VM can be pinged.
--compute-health	Performs a compute health check.
--general-health	Verifies ESXi entries across all sources, checks the Postgres DB operational status for hosts, checks ESXi for error dumps, and gets NSX Manager and cluster status.

Table continued on next page

Continued from previous page

Option	Description
--get-host-ips	Returns server information.
--health-check	Performs all available health checks.
--ntp-health	Verifies whether the time on the components is synchronized with the NTP server in the VMware Cloud Builder appliance.
--services-health	Performs a services health check to confirm whether services are running
--run-vsan-checks	Runs proactive vSAN tests to verify the ability to create VMs within the vSAN disks.

Sample Output

The following text is a sample output from an --ntp-health operation.

```
root@cloud-builder [ /opt/vmware/sddc-support ]# ./sos --ntp-health --skip-known-host --force
Welcome to Supportability and Serviceability(SoS) utility!

User passed --force flag, Running SOS from Cloud Builder VM, although Bringup is completed and SDDC Manager is available. Please expect failures with SoS operations.

Health Check : /var/log/vmware/vcf/sddc-support/healthcheck-2020-02-11-23-03-53-24681
Health Check log : /var/log/vmware/vcf/sddc-support/healthcheck-2020-02-11-23-03-53-24681/sos.log

SDDC Manager : sddc-manager.vrack.vsphere.local

NTP : GREEN

+-----+-----+-----+-----+
| SL# |           Area           | Title | State |
+-----+-----+-----+-----+
| 1 | ESXi : esxi-1.vrack.vsphere.local | ESX Time | GREEN |
| 2 | ESXi : esxi-2.vrack.vsphere.local | ESX Time | GREEN |
| 3 | ESXi : esxi-3.vrack.vsphere.local | ESX Time | GREEN |
| 4 | ESXi : esxi-4.vrack.vsphere.local | ESX Time | GREEN |
| 5 | vCenter : vcenter-1.vrack.vsphere.local | NTP Status | GREEN |
+-----+-----+-----+-----+
```

Legend:

GREEN - No attention required, health status is NORMAL

YELLOW - May require attention, health status is WARNING

RED - Requires immediate attention, health status is CRITICAL

Health Check completed successfully for : [NTP-CHECK]

The following text is sample output from a `--vm-screenshots` log collection operation.

```
root@cloud-builder [ /opt/vmware/sddc-support ]# ./sos --vm-screenshots
--skip-known-host --force
```

Welcome to Supportability and Serviceability(SoS) utility!

User passed `--force` flag, Running SOS from Cloud Builder VM, although Bringup is completed and SDDC Manager is available. Please expect failures with SoS operations.

Logs : /var/log/vmware/vcf/sddc-support/sos-2018-08-24-10-50-20-8013

Log file : /var/log/vmware/vcf/sddc-support/sos-2018-08-24-10-50-20-8013/sos.log

Log Collection completed successfully for : [VMS_SCREENSHOT]

VMware Cloud Builder Log Files

VMware Cloud Builder contains various log files for different components of the system.

VMware Cloud Builder has a number of components which are used during the bring-up process, each component generates a log file which can be used for the purpose of troubleshooting. The components and their purpose are:

- **JsonGenerator:** Used to convert the deployment parameter workbook into the required configuration file (JSON) that is used by the Bringup Validation Service and Bringup Service.
- **Bringup Service:** Used to perform the validation of the configuration file (JSON), the ESXi hosts and infrastructure where VMware Cloud Foundation will be deployed, and to perform the deployment and configuration of the management domain components and the first cluster.
- **Supportability and Serviceability (SoS) Utility:** A command line utility for troubleshooting deployment issues.

The following table describes the log file locations:

Component	Log Name	Location
JsonGenerator	<code>jsongenerator-timestamp</code>	<code>/var/log/vmware/vcf/sddc-support/</code>
Bringup Service	<code>vcf-bringup.log</code>	<code>/var/log/vmware/vcf/bringup/</code>
	<code>vcf-bringup-debug.log</code>	<code>/var/log/vmware/vcf/bringup/</code>
	<code>rest-api-debug.log</code>	<code>/var/log/vmware/vcf/bringup/</code>
SoS Utility	<code>sos.log</code>	<code>/var/log/vmware/vcf/sddc-support/sos-timestamp/</code>

Getting Started with SDDC Manager

You use SDDC Manager to perform administration tasks on your VMware Cloud Foundation instance. The SDDC Manager UI provides an integrated view of the physical and virtual infrastructure and centralized access to manage the physical and logical resources.

You work with the SDDC Manager UI by loading it in a web browser. For the list of supported browsers and versions, see the *Release Notes*.

Log in to the SDDC Manager User Interface

Connect to the SDDC Manager appliance by logging into the SDDC Manager UI using a supported web browser.

To log in, you need the SDDC Manager IP address or FQDN and the password for the single-sign on user (for example `administrator@vsphere.local`). You added this information to the deployment parameter workbook before bring-up.

1. In a web browser, type one of the following.
 - `https://FQDN` where *FQDN* is the fully-qualified domain name of the SDDC Manager appliance.
 - `https://IP_address` where *IP_address* is the IP address of the SDDC Manager appliance.
2. Log in to the SDDC Manager UI with vCenter Server Single Sign-On user credentials.

You are logged in to SDDC Manager UI and the Dashboard page appears in the web browser.

Guided SDDC Manager Onboarding

VMware Cloud Foundation includes an onboarding dashboard to help you with configuring a healthy SDDC Manager environment.

This dashboard appears when you log into SDDC Manager. It provides a walk-through for initial configuration, including the recommended order for completing each task. After completing the walk-through, a banner at the top of the screen offers a tour of the SDDC Manager UI.

You can skip sections and exit out of the guided setup at any point. This dashboard automatically shows unless you click "Don't show onboarding screen again" and close the page. Clicking this option also prevents the optional guided tour from automatically displaying in the future.

Use the Help Icon in the upper-right corner of the page to later access the onboarding dashboard and guided tour.

Tour of the SDDC Manager User Interface

The SDDC Manager UI provides a single point of control for managing and monitoring your VMware Cloud Foundation instance and for provisioning workload domains.

You use the navigation bar to move between the main areas of the user interface.

Navigation Bar

The navigation bar is available on the left side of the interface and provides a hierarchy for navigating to the corresponding pages.

Category	Functional Areas
Dashboard	<p>The Dashboard provides the high-level administrative view for SDDC Manager in the form of widgets. There are widgets for Solutions; Workload Domains; Host Types and Usage; Ongoing and Scheduled Updates; Update History; CPU, Memory, Storage Usage; and Recent Tasks.</p> <p>You can control the widgets that are displayed and how they are arranged on the dashboard.</p> <ul style="list-style-type: none"> • To rearrange widgets, click the heading of the widget and drag it to the desired position. • To hide a widget, hover the mouse anywhere over the widget to reveal the X in the upper-right corner, and click the X. • To add a widget, click the three dots in the upper right corner of the page and select Add New Widgets. This displays all hidden widgets. Select a widget and click Add.
Solutions	<p>Solutions include the following section:</p> <ul style="list-style-type: none"> • Kubernetes - Workload Management allows you to start a Workload Management deployment and view Workload Management cluster details.
Inventory	<p>Inventory includes the following sections:</p> <ul style="list-style-type: none"> • Workload Domains takes you to the Workload Domains page, which displays and provides access to all workload domains. <p>This page includes summary information about all workload domains, including domain type, storage usage, configuration status, owner, clusters, hosts and update availability. It also displays CPU, memory, and storage utilization for each workload domain, and collectively across all domains.</p> <ul style="list-style-type: none"> • Hosts takes you to the Hosts page, which displays and provides access to current hosts and controls for managing hosts. <p>This page includes detailed information about all hosts, including FQDN, host IP, network pool, configuration status, host state, cluster, and storage type. It also displays CPU and memory utilization for each host, and collectively across all hosts.</p>
Lifecycle Management	<p>Lifecycle Management includes the following sections:</p> <ul style="list-style-type: none"> • Release Versions displays the versions in your environment and the associated component versions in that release. • Bundle Management displays the available install, update, and upgrade bundles for your environment, and your bundle download history.

Table continued on next page

Continued from previous page

Category	Functional Areas
	<p>NOTE To access bundles, you must be logged in to your Broadcom Support Portal account through the Administration > Depot Settings page.</p>
Administration	<p>Administration includes the following sections:</p> <ul style="list-style-type: none"> • Network Settings allows you to update the DNS and NTP servers that VMware Cloud Foundation uses. • Licensing allows you to manage VMware product licenses. You can also add licenses for the component products in your VMware Cloud Foundation deployment. • Single Sign On allows you to manage VMware Cloud Foundation users and groups, including adding users and groups and assigning roles. You can also configure identity providers for VMware Cloud Foundation. • Proxy Settings allows you to configure a proxy server to download install and upgrade bundles from the VMware Depot. • Depot Settings allows you to log in to your Broadcom Support Portal and Dell accounts to download install and upgrade bundles. • VMware Aria Suite allows you to deploy VMware Aria Suite Lifecycle and configure connections between workload domains and VMware Aria Suite products. • Backup allows you to register an external SFTP server with SDDC Manager for backing up SDDC Manager and NSX Managers. You can also configure the backup schedule for SDDC Manager. • VMware CEIP to join or leave the VMware Customer Experience Improvement Program.
Security	<ul style="list-style-type: none"> • Password Management allows password management actions, such as rotation, updates and remediation. • Certificate Authority allows you to integrate with your Microsoft Certificate Authority Server.
Developer Center	<p>The VMware Cloud Foundation Developer Center includes the following sections:</p> <ul style="list-style-type: none"> • Overview: API reference documentation. Includes information and steps for all the Public APIs supported by VMware Cloud Foundation. • API Explorer: Lists the APIs and allows you to invoke them directly on your VMware Cloud Foundation system.

Log out of the SDDC Manager User Interface

Log out of the SDDC Manager UI when you have completed your tasks.

1. In the SDDC Manager UI, click the logged-in account name in the upper right corner.

2. Click **Log out**.


Configure the Customer Experience Improvement Program Settings for VMware Cloud Foundation

VMware Cloud Foundation participates in the VMware Customer Experience Improvement Program (CEIP). You can choose to activate or deactivate CEIP for your VMware Cloud Foundation instance.

The Customer Experience Improvement Program provides VMware with information that allows VMware to improve its products and services, to fix problems, and to advise you on how best to deploy and use our products. As part of the CEIP, VMware collects technical information about your organization's use of the VMware products and services regularly in association with your organization's VMware license keys. This information does not personally identify any individual. For additional information regarding the CEIP, refer to the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>.

You can activate or deactivate CEIP across all the components deployed in VMware Cloud Foundation by the following methods:

- When you log into SDDC Manager for the first time, a pop-up window appears. The **Join the VMware Customer Experience Program** option is selected by default. Deselect this option if you do not want to join CEIP. Click **Apply**.



- You can activate or deactivate CEIP from the Administration tab in the SDDC Manager UI.

1. In the navigation pane, click **Administration** > **VMware CEIP**.
2. To activate CEIP, select the **Join the VMware Customer Experience Improvement Program** option.
3. To deactivate CEIP, deselect the **Join the VMware Customer Experience Improvement Program** option.

Managing Certificates in VMware Cloud Foundation

You can use the SDDC Manager UI to manage certificates in a VMware Cloud Foundation instance, including integrating a certificate authority, generating and submitting certificate signing requests (CSR) to a certificate authority, and downloading and installing certificates.

Starting with VMware Cloud Foundation 5.2.1, you can also manage certificates using the vSphere Client.

This section provides instructions for the SDDC Manager UI to:

- Use OpenSSL as a certificate authority, which is a native option in SDDC Manager.
- Integrate with Microsoft Active Directory Certificate Services.
- Provide signed certificates from another external Certificate Authority.

You can manage the certificates for the following components.

- vCenter Server
- NSX Manager
- VMware Avi Load Balancer (formerly known as NSX Advanced Load Balancer)
- SDDC Manager
- VxRail Manager
- VMware Aria Suite Lifecycle

NOTE

Use VMware Aria Suite Lifecycle to manage certificates for the other VMware Aria Suite components.

You replace certificates for the following reasons:

- A certificate has expired or is nearing its expiration date.
- A certificate has been revoked by the issuing certificate authority.
- You do not want to use the default VMCA-signed certificates.
- Optionally, when you create a new workload domain.

It is recommended that you replace all certificates after completing the deployment of the VMware Cloud Foundation management domain. After you create a new VI workload domain, you can replace certificates for the appropriate components as needed.

View Certificate Information

You can view details of an applied certificate for a resource directly through the SDDC Manager UI.

The SDDC Manager UI provides a banner notification for any certificates that are expiring in the next 30 days.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the **Workload Domains** page, from the table, in the domain column click the domain you want to view.
3. On the domain summary page, click the **Certificates** tab.

Summary Services Updates Update History Hosts Clusters Edge Clusters Certificates								
GENERATE CSRS			GENERATE SIGNED CERTIFICATES			INSTALL CERTIFICATES		
						DOWNLOAD CSR		UPLOAD AND INSTALL CERTIFICATES
<input type="checkbox"/>	Resource Type	Issuer	Issued To	Valid From	Valid Until	Status	Certificate Operation Status	
<input type="checkbox"/>	> vcenter	CA	vcenter-vsan.vrack.vsphere.io...	Jun 8, 2023	Jun 7, 2025	Active	CSR Generation - NOT STARTED	
<input type="checkbox"/>	> nsx	CA	nsxt-manager-1-cls1.vrack.vsphere.local	Jun 8, 2023	Sep 10, 2025	Active	CSR Generation - NOT STARTED	
<input type="checkbox"/>	> nsx	CA	vip-nsxmanager-cls1.vrack.vsphere.local	Jun 8, 2023	Sep 10, 2025	Active	CSR Generation - NOT STARTED	

This tab lists the certificates for each resource type associated with the workload domain. It displays the following details:

- Resource type
- Issuer, the certificate authority name
- Resource hostname
- Valid From
- Valid Until
- Certificate status: Active, Expiring, or Expired.
- Certificate operation status

4. To view certificate details, expand the resource next to the Resource Type column.

Configure VMware Cloud Foundation to Use Microsoft CA-Signed Certificates

VMware Cloud Foundation supports the ability to manage certificates by integrating with Microsoft Active Directory Certificate Services (Microsoft CA). Before you can perform certificate operations using the SDDC Manager UI you must ensure that the Microsoft Certificate Authority is configured correctly.

Complete the below tasks to manage Microsoft CA-Signed certificates using SDDC Manager.

Prepare Your Microsoft Certificate Authority to Allow SDDC Manager to Manage Certificates

To ensure secure and operational connectivity between the SDDC components, you apply signed certificates provided by a Microsoft Certificate Authority for the SDDC components.

You use SDDC Manager to generate the certificate signing request (CSRs) and request a signed certificate from the Microsoft Certificate Authority. SDDC Manager is then used to install the signed certificates to SDDC components it manages. In order to achieve this the Microsoft Certificate Authority must be configured to allow integration with SDDC Manager.

Install Microsoft Certificate Authority Roles

Install the Certificate Authority and Certificate Authority Web Enrollment roles on the Microsoft Certificate Authority server to facilitate certificate generation from SDDC Manager.

NOTE

When connecting SDDC Manager to Microsoft Active Directory Certificate Services, ensure that Web Enrollment role is installed on the same machine where the Certificate Authority role is installed. SDDC Manager can't request and sign certificates automatically if the two roles (Certificate Authority and Web Enrollment roles) are installed on different machines.

1. Log in to the Microsoft Certificate Authority server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Add roles to Microsoft Certificate Authority server.
 - a) Click **Start > Run**, enter `ServerManager`, and click **OK**.
 - b) From the **Dashboard**, click **Add roles and features** to start the **Add Roles and Features** wizard.
 - c) On the **Before you begin** page, click **Next**.
 - d) On the **Select installation type** page, click **Next**.
 - e) On the **Select destination server** page, click **Next**.
 - f) On the **Select server roles** page, under **Active Directory Certificate Services**, select **Certification Authority** and **Certification Authority Web Enrollment** and click **Next**.
 - g) On the **Select features** page, click **Next**.
 - h) On the **Confirm installation selections** page, click **Install**.

Configure the Microsoft Certificate Authority for Basic Authentication

Configure the Microsoft Certificate Authority with basic authentication to allow SDDC Manager the ability to manage signed certificates.

The Microsoft Certificate Authority and IIS must be installed on the same server.

1. Log in to the Active Directory server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Add Basic Authentication to the Web Server (IIS).
 - a) Click **Start > Run**, enter `ServerManager`, and click **OK**.
 - b) From the **Dashboard**, click **Add roles and features** to start the **Add Roles and Features** wizard.
 - c) On the **Before you begin** page, click **Next**.
 - d) On the **Select installation type** page, click **Next**.
 - e) On the **Select destination server** page, click **Next**.
 - f) On the **Select server roles** page, under **Web Server (IIS) > Web Server > Security**, select **Basic Authentication** and click **Next**.
 - g) On the **Select features** page, click **Next**.
 - h) On the **Confirm installation selections** page, click **Install**.
3. Configure the certificate service template and CertSrv web site, for basic authentication.
 - a) Click **Start > Run**, enter `Inetmgr.exe` and click **OK** to open the **Internet Information Services Application Server Manager**.
 - b) Navigate to *your_server* > **Sites > Default Web Site > CertSrv**.
 - c) Under **IIS**, double-click **Authentication**.
 - d) On the **Authentication** page, right-click **Basic Authentication** and click **Enable**.
 - e) In the navigation pane, select **Default Web Site**.
 - f) In the **Actions** pane, under **Manage Website**, click **Restart** for the changes to take effect.

Create and Add a Microsoft Certificate Authority Template

You must set up a certificate template in the Microsoft Certificate Authority. The template contains the certificate authority attributes for signing certificates for the VMware Cloud Foundation components. After you create the template, you add it to the certificate templates of the Microsoft Certificate Authority.

1. Log in to the Active Directory server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Click **Start > Run**, enter `certtmpl.msc`, and click **OK**.
3. In the **Certificate Template Console** window, under **Template Display Name**, right-click **Web Server** and select **Duplicate Template**.
4. In the **Properties of New Template** dialog box, click the **Compatibility** tab and configure the following values.

Setting	Value
Certification Authority	Windows Server 2008 R2
Certificate recipient	Windows 7 / Server 2008 R2

5. In the **Properties of New Template** dialog box, click the **General** tab and enter a name for example, `VMware` in the **Template display name** text box.
6. In the **Properties of New Template** dialog box, click the **Extensions** tab and configure the following.
 - a) Click **Application Policies** and click **Edit**.
 - b) Click **Server Authentication**, click **Remove**, and click **OK**.
 - c) Click **Basic Constraints** and click **Edit**.
 - d) Click the **Enable this extension** check box and click **OK**.
 - e) Click **Key Usage** and click **Edit**.
 - f) Click the **Signature is proof of origin (nonrepudiation)** check box, leave the defaults for all other options and click **OK**.
7. In the **Properties of New Template** dialog box, click the **Subject Name** tab, ensure that the **Supply in the request** option is selected, and click **OK** to save the template.
8. Add the new template to the certificate templates of the Microsoft CA.
 - a) Click **Start > Run**, enter `certsrv.msc`, and click **OK**
 - b) In the **Certification Authority** window, expand the left pane, right-click **Certificate Templates**, and select **New > Certificate Template to Issue**.
 - c) In the **Enable Certificate Templates** dialog box, select **VMware**, and click **OK**.

Assign Certificate Management Privileges to the SDDC Manager Service Account

Before you can use the Microsoft Certificate Authority and the pre-configured template, it is recommended to configure least privilege access to the Microsoft Active Directory Certificate Services using an Active Directory user account as a restricted service account.

- Create a user account in Active Directory with Domain Users membership. For example, `svc-vcf-ca`.
1. Log in to the Microsoft Certificate Authority server by using a Remote Desktop Protocol (RDP) client.

FQDN	<i>Active Directory Host</i>
User	Active Directory administrator
Password	<i>ad_admin_password</i>

2. Configure least privilege access for a user account on the Microsoft Certificate Authority.
 - a) Click **Start > Run**, enter `certsrv.msc`, and click **OK**.
 - b) Right-click the certificate authority server and click **Properties**.
 - c) Click the **Security** tab, and click **Add**.
 - d) Enter the name of the user account and click **OK**.
 - e) In the **Permissions for** section configure the permissions and click **OK**.

Setting	Value (Allow)
Read	Deselected
Issue and Manage Certificates	Selected
Manage CA	Deselected
Request Certificates	Selected

3. Configure least privilege access for the user account on the Microsoft Certificate Authority Template.
 - a) Click **Start** > **Run**, enter `certtmpl.msc`, and click **OK**.
 - b) Right-click the VMware template and click **Properties**.
 - c) Click the **Security** tab, and click **Add**.
 - d) Enter the `svc-vcf-ca` service account and click **OK**.
 - e) In the **Permissions for** section configure the permissions and click **OK**.

Setting	Value (Allow)
Full Control	Deselected
Read	Selected
Write	Deselected
Enroll	Selected
Autoenroll	Deselected

Configure a Microsoft Certificate Authority in SDDC Manager

You configure a connection between SDDC Manager and a Microsoft Certificate Authority by entering your service account credentials.

- Verify connectivity between SDDC Manager and the Microsoft Certificate Authority Server. See [VMware Ports and Protocols](#).
- Verify that the Microsoft Certificate Authority Server has the correct roles installed on the same machine where the Certificate Authority role is installed. See [Install Microsoft Certificate Authority Roles](#).
- Verify the Microsoft Certificate Authority Server has been configured for basic authentication. See [Configure the Microsoft Certificate Authority for Basic Authentication](#).
- Verify a valid certificate template has been configured on the Microsoft Certificate Authority. See [Create and Add a Microsoft Certificate Authority Template](#).
- Verify least privileged user account has been configured on the Microsoft Certificate Authority Server and Template. See [Assign Certificate Management Privileges to the SDDC Manager Service Account](#).
- Verify that time is synchronized between the Microsoft Certificate Authority and the SDDC Manager appliance. Each system can be configured with a different timezone, but it is recommended that they receive their time from the same NTP source.

1. In the navigation pane, click **Security** > **Certificate Authority**.
2. Click **Edit**.

Configure Certificate Authority EDIT

Certificate Authority Type

CA Server URL

User Name

Password

Template Name

SAVE

- Configure the settings and click **Save**.

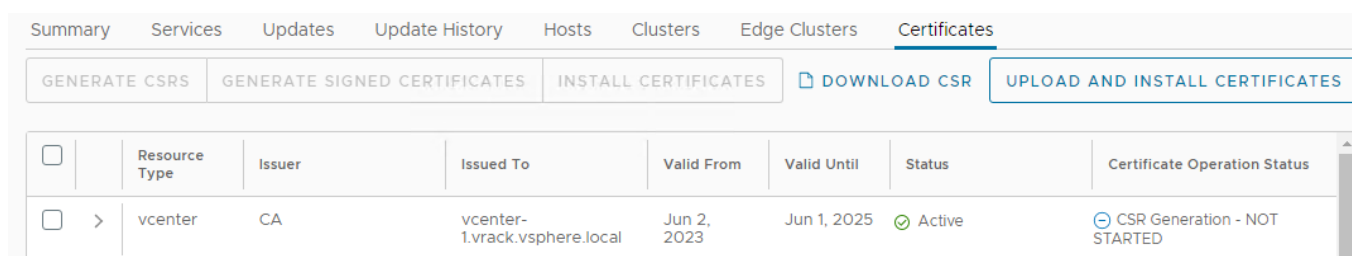
Setting	Value
Certificate Authority Type	Microsoft
CA Server URL	Specify the URL for the issuing certificate authority. This address must begin with <code>https://</code> and end with <code>certsrv</code> . For example, <code>https://ca.rainpole.io/certsrv</code> .
User Name	Enter a least privileged service account. For example, <code>svc-vcf-ca</code> .
Password	Enter the password for the least privileged service account.
Template Name	Enter the issuing certificate template name. You must create this template in Microsoft Certificate Authority. For example, VMware.

- In the **CA Server Certificate Details** dialog box, click **Accept**.

Install Microsoft CA-Signed Certificates using SDDC Manager

Replace the self-signed certificates with signed certificates from the Microsoft Certificate Authority by using SDDC Manager.

- In the navigation pane, click **Inventory** > **Workload Domains**.
- On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
- On the domain summary page, click the **Certificates** tab.



4. Generate CSR files for the target components.
 - a) From the table, select the check box for the resource type for which you want to generate a CSR.
 - b) Click **Generate CSRs**.
 - c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.
 - e) On the **Summary** dialog, click **Generate CSRs**.
5. Generate signed certificates for each component.
 - a) From the table, select the check box for the resource type for which you want to generate a signed certificate for.
 - b) Click **Generate Signed Certificates**.
 - c) In the **Generate Certificates** dialog box, from the **Select Certificate Authority** drop-down menu, select **Microsoft**.
 - d) Click **Generate Certificates**.
6. Install the generated signed certificates for each component.
 - a) From the table, select the check box for the resource type for which you want to install a signed certificate.
 - b) Click **Install Certificates**.

Configure VMware Cloud Foundation to Use OpenSSL CA-Signed Certificates

VMware Cloud Foundation supports the ability to manage certificates using OpenSSL configured on the SDDC Manager appliance.

Complete the following tasks to be able to manage OpenSSL-signed certificates issued by SDDC Manager.

Configure OpenSSL-signed Certificates in SDDC Manager

To generate OpenSSL-signed certificates for the VMware Cloud Foundation components you must first configure the certificate authority details.

1. In the navigation pane, click **Security** > **Certificate Authority**.
2. Click **Edit**.
3. Configure the settings and click **Save**.

Setting	Value
Certificate Authority	OpenSSL
Common Name	Specify the FQDN of the SDDC Manager appliance.
Organizational Unit	Use this field to differentiate between the divisions within your organization with which this certificate is associated.

Table continued on next page


Continued from previous page

Setting	Value
Organization	Specify the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Specify the city or the locality where your company is legally registered.
State	Enter the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Select the country where your company is registered. This value must use the ISO 3166 country code.

Install OpenSSL-signed Certificates using SDDC Manager

Replace the self-signed certificates with OpenSSL-signed certificates generated by SDDC Manager.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
3. On the domain summary page, click the **Certificates** tab.



4. Generate CSR files for the target components.
 - a) From the table, select the check box for the resource type for which you want to generate a CSR.
 - b) Click **Generate CSRs**.
The **Generate CSRs** wizard opens.
 - c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal

Table continued on next page

Continued from previous page

Option	Description
	registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.

You can enter multiple values separated by comma (,), semicolon (;), or space (). For NSX, you can enter the subject alternative name for each node along with the Virtual IP (primary) node.

NOTE

Wildcard subject alternate name, such as *.example.com is not recommended.

- e) On the **Summary** dialog, click **Generate CSRs**.
5. Generate signed certificates for each component.
- From the table, select the check box for the resource type for which you want to generate a signed certificate.
 - Click **Generate Signed Certificates**.
 - In the **Generate Certificates** dialog box, from the **Select Certificate Authority** drop-down menu, select **OpenSSL**.
 - Click **Generate Certificates**.
6. Install the generated signed certificates for each component.
- From the table, select the check box for the resource type for which you want to install a signed certificate.
 - Click **Install Certificates**.

Install Third-Party CA-Signed Certificates Using Server Certificate and Certificate Authority Files

VMware Cloud Foundation supports two ways to install third-party certificates. This procedure describes the new method, which is the default method for VMware Cloud Foundation 4.5.1 and later.

If you prefer to use the legacy method for installing third-party CA-signed certificates, see [Install Third-Party CA-Signed Certificates in VMware Cloud Foundation Using a Certificate Bundle](#).

- In the navigation pane, click **Inventory** > **Workload Domains**.
- On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
- On the domain summary page, click the **Certificates** tab.
- Generate CSR files for the target components.
 - From the table, select the check box for the resource type for which you want to generate a CSR.
 - Click **Generate CSRs**.

The **Generate CSRs** wizard opens.

- c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.

You can enter multiple values separated by comma (,), semicolon (;), or space (). For NSX, you can enter the subject alternative name for each node along with the Virtual IP (primary) node.

NOTE

Wildcard subject alternative name, such as *.example.com are not recommended.

- e) On the **Summary** dialog, click **Generate CSRs**.

- Download and save the CSR files by clicking **Download CSR**.
- When the downloads complete, request signed certificates from your third-party Certificate Authority for each .csr.
- After you receive the signed certificates, open the SDDC Manager UI and click **Upload and Install**.
- In the **Install Signed Certificates** dialog box, select the resource for which you want to install a signed certificate. The drop-down menu includes all resources for which you have generated and downloaded CSRs.
- Select a **Source** and enter the required information.

Source	Required Information
Paste Text	Copy and paste the: <ul style="list-style-type: none"> • Server Certificate • Certificate Authority Paste the server certificate and the certificate authority in PEM format (base64-encoded) . For example: <pre>-----BEGIN CERTIFICATE----- <certificate content></pre>

Table continued on next page

Continued from previous page

Source	Required Information
	<pre>-----END CERTIFICATE----- If the Certificate Authority includes intermediate certificates, it should be in the following format: -----BEGIN CERTIFICATE----- <Intermediate certificate content> -----END CERTIFICATE----- -----BEGIN CERTIFICATE----- <Root certificate content> -----END CERTIFICATE-----</pre>
File Upload	Click Browse to upload the: <ul style="list-style-type: none"> • Server Certificate • Certificate Authority Files with .crt, .cer, .pem, .p7b and .p7c extensions are supported.
Certificate Chain	Click Browse to upload the certificate chain. Files with .crt, .cer, .pem, .p7b and .p7c extensions are supported.

10. Click **Validate**.

If validation fails, resolve the issues and try again, or click **Remove** to skip the certificate installation.

11. To install a signed certificate for another resource, click **Add Another** and repeat steps 8-10 for each resource.


12. Once all signed certificates have been validated successfully, click **Install**.

Install Third-Party CA-Signed Certificates in VMware Cloud Foundation Using a Certificate Bundle

VMware Cloud Foundation supports two ways to install third-party certificates. This procedure describes the legacy method of using a certificate bundle. To use the legacy method, you must modify your preferences and then use this procedure to generate CSRs, sign the CSRs with a third-party CA, and finally upload and install the certificates.

VMware Cloud Foundation 4.5.1 introduces a new method for installing third-party CA-signed certificates. By default, VMware Cloud Foundation use the new method. See [Install Third-Party CA-Signed Certificates Using Server Certificate and Certificate Authority Files](#) for information using the new method. If you prefer to use the legacy method, you must modify your preferences.

1. In the SDDC Manager UI, click the logged in user and select **Preferences**.



2. Use the toggle to switch to legacy certificate management.

Revert to Legacy Certificate Management



Uploading CA-signed certificates from a third-party Certificate Authority using the legacy method requires that you collect the relevant certificate files in the correct format and then create a single `.tar.gz` file with the contents. It's important that you create the correct directory structure within the `.tar.gz` file as follows:

- The name of the top-level directory must exactly match the name of the workload domain as it appears in the list on the **Inventory > Workload Domains**. For example, `sfo-m01`.
 - The PEM-encoded root CA certificate chain file (must be named `rootca.crt`) must reside inside this top-level directory. The `rootca.crt` chain file contains a root certificate authority and can have `n` number of intermediate certificates.

For example:

```
-----BEGIN CERTIFICATE-----
<Intermediate1 certificate content>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Intermediate2 certificate content>
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
<Root certificate content>
-----END CERTIFICATE-----
```

In the above example, there are two intermediate certificates, `intermediate1` and `intermediate2`, and a root certificate. `intermediate1` must use the certificate issued by `intermediate2` and `intermediate2` must use the certificate issued by Root CA.

- The root CA certificate chain file, intermediate certificates, and root certificate must contain the `Basic Constraints` field with value `CA:TRUE`.
- This directory must contain one sub-directory for each component resource for which you want to replace the certificates.
- Each sub-directory must exactly match the resource hostname of a corresponding component as it appears in the Resource Hostname column in the **Inventory > Workload Domains > Certificates** tab.

For example, `nsxManager.vrack.vsphere.local`, `vcenter-1.vrack.vsphere.local`, and so on.

- Each sub-directory must contain the corresponding `.csr` file, whose name must exactly match the resource as it appears in the Resource Hostname column in the **Inventory › Workload Domains › Certificates** tab.
- Each sub-directory must contain a corresponding `.crt` file, whose name must exactly match the resource as it appears in the Resource Hostname column in the **Inventory › Workload Domains › Certificates** tab. The content of the `.crt` files must end with a newline character.

For example, the `nsxManager.vrack.vsphere.local` sub-directory would contain the `nsxManager.vrack.vsphere.local.crt` file.

- All certificates including `rootca.crt` must be in UNIX file format.
- Additional requirements for NSX certificates:
 - Server certificate (`NSX_FQDN.crt`) must contain the `Basic Constraints` field with value `CA:FALSE`.
 - If the NSX certificate contains HTTP or HTTPS based CRL Distribution Point it must be reachable from the server.
 - The extended key usage (EKU) of the generated certificate must contain the EKU of the CSR generated.

NOTE

All resource and hostname values can be found in the list on the **Inventory › Workload Domains › Certificates** tab.

1. In the navigation pane, click **Inventory › Workload Domains**.
2. On the **Workload Domains** page, from the table, in the domain column click the workload domain you want to view.
3. On the domain summary page, click the **Certificates** tab.
4. Generate CSR files for the target components.
 - a) From the table, select the check box for the resource type for which you want to generate a CSR.
 - b) Click **Generate CSRs**.
The **Generate CSRs** wizard opens.
 - c) On the **Details** dialog, configure the settings and click **Next**.

Option	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.

Table continued on next page

Continued from previous page

Option	Description
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d) (Optional) On the **Subject Alternative Name** dialog, enter the subject alternative name(s) and click **Next**.

You can enter multiple values separated by comma (,), semicolon (;), or space (). For NSX, you can enter the subject alternative name for each node along with the Virtual IP (primary) node.

NOTE

Wildcard subject alternative name, such as *.example.com are not recommended.

- e) On the **Summary** dialog, click **Generate CSRs**.
5. Download and save the CSR files to the directory by clicking **Download CSR**.
 6. Complete the following tasks outside of the SDDC Manager UI:
 - a) Verify that the different .csr files have successfully generated and are allocated in the required directory structure.
 - b) Request signed certificates from a Third-party Certificate authority for each .csr.
 - c) Verify that the newly acquired .crt files are correctly named and allocated in the required directory structure.
 - d) Create a new .tar.gz file of the directory structure ready for upload to SDDC Manager. For example: `<domain name>.tar.gz`.
 7. Click **Upload and Install**.
 8. In the **Upload and Install Certificates** dialog box, click **Browse** to locate and select the newly created `<domain name>.tar.gz` file and click **Open**.
 9. Click **Upload**.
 10. If the upload is successful, click **Install Certificate**. The Certificates tab displays a status of Certificate Installation is in progress.

Remove Old or Unused Certificates from SDDC Manager

Old or unused certificates are stored in a trust store in SDDC Manager. You can delete old certificates using the VMware Cloud Foundation API.

See [Delete Trusted Certificate](#) in the *VMware Cloud Foundation API Reference Guide* for more information.

1. Log in to the SDDC Manager UI as a user with the ADMIN role.
For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).
2. In the navigation pane, click **Developer Center > API Explorer**.
3. Browse to and expand **API Categories > Trusted Certificates**.
4. Expand **GET /v1/sddc-manager/trusted-certificates** and click **EXECUTE**.
5. In the Response, click `TrustedCertificate` and copy the alias for the certificate you want to remove.
6. Expand **DELETE /v1/sddc-manager/trusted-certificates/{alias}**, enter the alias, and click **EXECUTE**.

Try it out

Parameter	Value	Type	Description / Data Type
alias (required)	<input type="text"/>	path	Certificate Alias Data Type: string

EXECUTE

COPY RESPONSE

DOWNLOAD

Managing License Keys in VMware Cloud Foundation

You can add component license keys in the SDDC Manager UI or add a solution license key in vSphere Client.

Starting with VMware Cloud Foundation 5.1.1, you can license VMware Cloud Foundation components using a solution license key or individual component license keys.

NOTE

VMware Cloud Foundation 5.1.1 supports a combination of solution and component license keys. For example, `Workload Domain 1` can use component license keys and `Workload Domain 2` can use the solution license key.

For more information about the VCF solution license key, VMware vSphere 8 Enterprise Plus for VCF, see <https://knowledge.broadcom.com/external/article?articleNumber=319282>.

SDDC Manager does not manage the solution license key. If you are using a solution license key, VMware Cloud Foundation components are deployed in evaluation mode and then you use the vSphere Client to add and assign the solution key. See [Managing vSphere Licenses](#) for information about using a solution license key for VMware ESXi and vCenter Server. If you are using a solution license key, you must also add a separate VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

NOTE

VMware vCenter Server, VMware NSX, VMware Aria Suite components, and VMware HCX are all licensed when you assign a solution license key to a vCenter Server.

Use the SDDC Manager UI to manage component license keys. If you entered component license keys in the deployment parameter workbook that you used to create the management domain, those component license keys appear in the Licensing screen of the SDDC Manager UI. You can add additional component license keys to support your requirements. You must have adequate license units available before you create a VI workload domain, add a host to a vSphere cluster, or add a vSphere cluster to a workload domain. Add the necessary component license keys before you begin any of these tasks.

Add a Component License Key in the SDDC Manager UI

You can use the SDDC Manager UI to add component license keys to the SDDC Manager inventory.

SDDC Manager does not manage solution license keys. See [Managing License Keys in VMware Cloud Foundation](#) for more information about solution license keys.

1. In the navigation pane, click **Administration > Licensing**.
2. Click **+ License Key**.



3. Select a product from the drop-down menu.
4. Enter the license key.

5. Enter a description for the license.
A description can help in identifying the license.

6. Click **Add**.

If you want to replace an existing license with a newly added license, you must add and assign the new license in the management UI (for example, vSphere Client or NSX Manager) of the component whose license you are replacing.

Edit a Component License Key Description in the SDDC Manager UI

If you have multiple component license keys for a product, the description can help in identifying the license key. For example, you may want to use one license key for high-performance workload domains and the other license key for regular workload domains.

1. In the navigation pane, click **Administration** > **Licensing**.
2. Click the vertical ellipsis (three dots) next to the license key and click **Edit Description**.

Edit Description

Remove

3. On the **Edit License Key Description** dialog, edit the description and click **Save**.

Delete a Component License Key in the SDDC Manager UI

Deleting a component license key removes it from the SDDC Manager inventory. If the license key has been applied to any workload domain, host, or vSphere cluster, it is not removed from them, but it cannot be applied to new workload domains, hosts, or vSphere clusters.

1. In the navigation pane, click **Administration** > **Licensing**.
2. Click the vertical ellipsis (three dots) next to the license key you want to delete and click **Remove**.

Edit Description

Remove

3. In the **Remove License key** dialog, click **Remove**.

The component license key is removed from the SDDC Manager inventory

Update Component License Keys for Workload Domain Components

You can use the SDDC Manager UI to update the license keys for components whose license keys have expired, are expiring, or are incompatible with upgraded components.

The new component license key(s) must already be added to the SDDC Manager inventory. See [Add a Component License Key in the SDDC Manager UI](#).

You can update component license keys for:

- vCenter Server
- VMware NSX
- VMware vSAN
- ESXi

Updates are specific to the selected workload domain. If you want to update component license keys for multiple workload domains, you must update each workload domain separately.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. Click a workload domain name in the **Domain** column.
3. Select **Actions** > **Update Licenses**.
4. Read the overview and click **Next**.
5. Select one or more products to update and click **Next**.

Update Licenses

> ✓ Overview Overview of updating licenses for a workload domain

▼ 2. Product Selection vCenter

Select product(s) within this workload domain that require the license to be updated.

Select products

<input type="checkbox"/>	Product	License key status ⓘ	Available compatible license keys
<input checked="" type="checkbox"/>	vCenter	✓ Active	Yes
<input type="checkbox"/>	NSX	✓ Active	Yes
<input type="checkbox"/>	vSAN	✓ Active	Yes
<input type="checkbox"/>	ESXi	✓ Active	Yes

NEXT

3. vCenter License Apply license to vCenter

4. Review Review applied licenses

6. Select a component license key for each product.
For VMware vSAN and ESXi, you must select the clusters that you want to update with new license keys.
7. Review the new component license keys and click **Submit**.

ESXi Lockdown Mode

You can activate or deactivate normal lockdown mode in VMware Cloud Foundation to increase the security of your ESXi hosts.

To activate or deactivate normal lockdown mode in VMware Cloud Foundation, you must perform operations through the vCenter Server. For information on how to activate or deactivate normal lockdown mode, see [vSphere Security](#).

You can activate normal lockdown mode on a host after the host is added to workload domain. VMware Cloud Foundation creates service accounts that can be used to access the hosts. Service accounts are added to the Exception Users list during the bring-up or host commissioning. You can rotate the passwords for the service accounts using the password management functionality in the SDDC Manager UI.

Managing Storage in VMware Cloud Foundation

To create and manage a workload domain, VMware Cloud Foundation requires at least one shared storage type for all ESXi hosts within a cluster. This initial shared storage type, known as principal storage, is configured during VxRail first run. Additional shared storage, known as supplemental storage, can be added using the vSphere Client after a cluster has been created.

Principal Storage

During the VxRail first run, you configure the initial shared storage type. This initial shared storage type is known as principal storage. Once created, the principal storage type for a cluster cannot be changed. However, a VI workload domain can include multiple clusters with unique principal storage types.

VMware Cloud Foundation supports the following types of principal storage:

- vSAN
 - vSAN Original Storage Architecture (vSAN OSA)
 - vSAN Express Storage Architecture (vSAN ESA)

NOTE

You cannot convert vSAN OSA to vSAN ESA or vice versa.

- VMFS on FC (Fibre Channel)

Supplemental Storage

Additional shared storage, known as supplemental storage, can be manually added or removed using the vSphere Client after a cluster has been created. All supplemental storage must be listed in the VMware Compatibility Guide. Multiple supplemental storage types can be presented to a cluster in the management domain or any VI workload domain.

VMware Cloud Foundation supports using the vSphere Client to add the following datastore types to a cluster:

- vSphere VMFS

vSAN Storage with VMware Cloud Foundation

vSAN is the preferred principal storage type for VMware Cloud Foundation. It is an enterprise-class storage integrated with vSphere and managed by a single platform. vSAN is optimized for flash storage and can non-disruptively expand capacity and performance by adding hosts to a cluster (scale-out) or by adding disks to a host (scale-up).

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	Yes	Yes	Yes
Supplemental	No	No	No

Prerequisites for vSAN Storage

To create a VI workload domain that uses vSAN as principal storage you must ensure the following:

- A minimum of three ESXi hosts that meet the vSAN hardware, cluster, software, networking and license requirements. For information, see the [vSAN Planning and Deployment Guide](#).
- Perform a VxRail first run specifying the vSAN configuration settings. For information on the VxRail first run, contact Dell Support.
- A valid vSAN license. See [Managing License Keys in VMware Cloud Foundation](#). You cannot use vSAN ESA without a qualifying license.

In some instances SDDC Manager may be unable to automatically mark the host disks as capacity. Follow the Mark Flash Devices as Capacity Using ESXCLI procedure in the [vSAN Planning and Deployment Guide](#).

Procedures for vSAN Storage

- To use vSAN as principal storage for a new VI workload domain, perform the VxRail first run and then create the VI workload domain. See [Creating VxRail VI Workload Domains](#).
- To use vSAN as principal storage for a new cluster, perform the VxRail first run and then add the VxRail cluster. See [Add a VxRail Cluster to a Workload Domain Using the SDDC Manager UI](#).

vSAN Original Storage Architecture (OSA)

With vSAN OSA, each host that contributes storage devices to the vSAN datastore must provide at least one device for flash cache and at least one device for capacity. The devices on the contributing host form one or more disk groups. Each disk group contains one flash cache device, and one or multiple capacity devices for persistent storage. Each host can be configured to use multiple disk groups.

vSAN OSA clusters can mount a remote datastore from other vSAN OSA clusters.

vSAN Express Storage Architecture (ESA)

With vSAN ESA, all storage devices claimed by vSAN contribute to capacity and performance. Each host's storage devices claimed by vSAN form a storage pool. The storage pool represents the amount of caching and capacity provided by the host to the vSAN datastore.

vSAN ESA clusters can mount a remote datastore from other vSAN ESA clusters.

To use vSAN ESA, you need:

- A direct or proxy internet connection OR a downloaded copy of the vSAN HCL JSON file

NOTE

SDDC Manager will keep the HCL file updated if it has direct or proxy internet connection.

- ESXi host disks to support vSAN ESA
- A vLCM image to manage clusters.

vSAN Compute Clusters

A vSAN compute cluster is a vSphere cluster with a small vSAN element that enables it to mount a remote datastore. The hosts in a compute cluster do not have local storage. A compute cluster can only mount a remote datastore from a cluster within the same workload domain.

A vSAN compute cluster can mount a datastore from one of the following cluster types:

- vSAN OSA
- vSAN ESA

Once you mount a remote datastore on a vSAN compute cluster, you can only mount additional datastores of the same cluster type.

NOTE: Datastores on clusters created outside of VMware Cloud Foundation cannot be mounted on VCF-created clusters. Likewise, clusters created outside of VMware Cloud Foundation cannot mount a datastore from a VCF-created cluster.

Fibre Channel Storage with VMware Cloud Foundation

Fibre Channel (FC) is a storage protocol that the SAN uses to transfer data traffic from ESXi hosts to shared storage. The protocol packages SCSI commands into FC frames. To connect to the FC SAN, the ESXi host uses Fibre Channel host bus adapters (HBAs).

Fibre Channel can be used as supplemental storage for the management domain and consolidated workload domains, however it can be used as principal storage for VI workload domains and can also be used a principal storage in a management domain converted from vSphere infrastructure.

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Principal	No	Only for a management domain converted from vSphere infrastructure	Yes

Table continued on next page

Continued from previous page

Storage Type	Consolidated Workload Domain	Management Domain	VI Workload Domain
Supplemental	Yes	Yes	Yes

Prerequisites for FC Storage

- A minimum of three ESXi hosts. Review the ESXi Fibre Channel SAN Requirements in the [vSphere Storage Guide](#).

NOTE

If you are using VMFS on FC as principal storage, and your VI workload domain is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- Perform a VxRail first run specifying the VMFS on FC configuration settings. For information on the VxRail first run, contact Dell Support.
- A pre-created VMFS datastore.

Procedures for FC Storage

- To use Fibre Channel as principal storage for a new VI workload domain, perform the VxRail first run and then create the VI workload domain. See [Creating VxRail VI Workload Domains](#).
- To use Fibre Channel as principal storage for a new cluster, perform the VxRail first run and then add the VxRail cluster. See [Add a VxRail Cluster to a Workload Domain Using the SDDC Manager UI](#)
- To use Fibre Channel as supplemental storage, see the [vSphere Storage Guide](#).

Sharing Remote Datastores with HCI Mesh for VI Workload Domains

HCI Mesh is a software-based approach for disaggregation of compute and storage resources in vSAN. HCI Mesh brings together multiple independent vSAN clusters by enabling cross-cluster utilization of remote datastore capacity within vCenter Server. HCI Mesh enables you to efficiently utilize and consume data center resources, which provides simple storage management at scale.

VMware Cloud Foundation supports sharing remote datastores with HCI Mesh for VI workload domains.

You can create HCI Mesh by mounting remote vSAN datastores on vSAN clusters and enable data sharing from the vCenter Server. It can take up to 5 minutes for the mounted remote vSAN datastores to appear in the .

It is recommended that you do not mount or configure remote vSAN datastores for vSAN clusters in the management domain.

For more information on sharing remote datastores with HCI Mesh, see [Sharing Remote Datastores with HCI Mesh](#).

NOTE

You cannot mount remote vSAN datastores on stretched clusters.

NOTE

After enabling HCI Mesh by mounting remote vSAN datastores, you can migrate VMs from the local datastore to a remote datastore. Since each cluster has its own VxRail Manager VM, you should not migrate VxRail Manager VMs to a remote datastore.

Managing Workload Domains in VMware Cloud Foundation

Workload domains are logical units that carve up the compute, network, and storage resources of the VMware Cloud Foundation system. The logical units are groups of ESXi hosts managed by vCenter Server instances with specific characteristics for redundancy and VMware best practices.

Each workload domain include these VMware capabilities by default:

- vCenter Server Appliance
- vSphere High Availability (HA)
- vSphere Distributed Resource Scheduler (DRS)
- vSphere Distributed Switch
- VMware vSAN
- NSX Manager Cluster

About VI Workload Domains

When deploying a workload domain, you specify the name, compute, and networking details for the VI workload domain. You then select the hosts for the VI workload domain and start the workflow.

When you deploy a new VI workload domain, VMware Cloud Foundation deploys a new vCenter Server for that workload domain. The vCenter Server is associated with a vCenter Single Sign-On Domain (SSO) to determine the local authentication space. Prior to VMware Cloud Foundation 5.0, the management vCenter Server and all VI workload domain vCenter Servers were members of a single vSphere SSO domain, joined together with vCenter Enhanced Linked Mode. Starting with VMware Cloud Foundation 5.0, when you deploy a new VI workload domain, you can choose to join the management domain SSO domain, or create a new SSO domain.

The workflow automatically:

- Deploys a vCenter Server Appliance for the new VI workload domain within the management domain. By using a separate vCenter Server instance per VI workload domain, software updates can be applied without impacting other VI workload domains. It also allows for each VI workload domain to have additional isolation as needed.
- Configures networking on each host.
- Configures vSAN storage on the ESXi hosts.
- For the first VI workload domain, the workflow deploys a cluster of three NSX Managers in the management domain and configures a virtual IP (VIP) address for the NSX Manager cluster. The workflow also configures an anti-affinity rule between the NSX Manager VMs to prevent them from being on the same host for high availability. Subsequent VI workload domains can share an existing NSX Manager cluster or deploy a new one. To share an NSX Manager cluster, the VI workload domains must use the same update method. The VI workload domains must both use vSphere Lifecycle Manager (vLCM) images, or they must both use vLCM baselines.
- By default, VI workload domains do not include any NSX Edge clusters and are isolated. To provide north-south routing and network services, add one or more NSX Edge clusters to a VI workload domain. See [Managing NSX Edge Clusters in VMware Cloud Foundation](#).

NOTE

Starting with VMware Cloud Foundation 5.2, when you deploy a new VI workload domain, it uses the same versions of vCenter Server and NSX Manager that the management domain uses. For example, if you applied an async patch to the vCenter Server in the management domain, a new VI workload domain will deploy the same patched version of vCenter Server.

Prerequisites for a Workload Domain

Review the prerequisites before you deploy a VI workload domain.

- If you plan to use DHCP for the NSX host overlay network, a DHCP server must be configured on the NSX host overlay VLAN for the VI workload domain. When VMware NSX creates NSX Edge tunnel endpoints (TEPs) for the VI workload domain, they are assigned IP addresses from the DHCP server.

NOTE

If you do not plan to use DHCP, you can use a static IP pool for the NSX host overlay network. The static IP pool is created or selected as part of VI workload domain creation.

- [Change the VxRail Manager IP Address](#)
- [Update the VxRail Manager Certificate](#)

- A minimum of three hosts available for the VI workload domain.

NOTE

If you are using VMFS on FC as principal storage, and the VI workload domain is using vSphere Lifecycle Manager images as the update method, then only two hosts are required. Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.

- The install bundles for the versions of NSX Manager and vCenter Server that are running in the management domain must be available in SDDC Manager before you can create a VI workload domain. For example, if you have patched the versions of NSX Manager and/or vCenter Server in the management domain to a version higher than what is listed in the BOM, you must download the new install bundles. You can refer to <https://knowledge.broadcom.com/external/article?legacyId=88287> for information about the install bundles required for specific async patches.
- Decide on a name for your VI workload domain. Each VI workload domain must have a unique name. It is good practice to include the region and site information in the name because resource object names (such as host and vCenter names) are generated based on the VI workload domain name. The name can be three to 20 characters long and can contain any combination of the following:
 - Lowercase alphabetic characters
 - Numbers

NOTE

Spaces are not allowed in any of the names you specify when creating a VI workload domain.

- Decide on the following passwords:
 - vCenter Server root password
 - NSX Manager admin password

Although the individual VMware Cloud Foundation components support different password requirements, you must set passwords following a common set of requirements across all components:

- Minimum length: 12
- Maximum length: 16
- At least one lowercase letter, one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ *
- Must NOT include:
 - A dictionary word
 - A palindrome
 - More than four monotonic character sequences
 - Three of the same consecutive characters
- Verify that you have the completed Planning and Preparation Workbook with the VI workload domain deployment option included.
- The IP addresses and Fully Qualified Domain Names (FQDNs) for the vCenter Server and NSX Manager instances must be resolvable by DNS.
- If you are using VMFS on FC storage for the VI workload domain, you must configure zoning, mount the associated volumes and create the datastore on the hosts.
- To use the **License Now** option, you must have valid license keys for the following products:
 - VMware NSX
 - vSAN (No license required for VMFS on FC)
 - vSphere

Because vSphere and vSAN licenses are per CPU, ensure that you have sufficient licenses for the ESXi hosts to be used for the VI workload domain. See [Managing License Keys in VMware Cloud Foundation](#).
- If you plan to deploy a VI workload domain that has its vSphere cluster at a remote location, you must meet the following requirements:
 - Dedicated WAN connectivity is required between central site and remote site.
 - Primary and secondary active WAN links are recommended for connectivity from the central site to the remote site. The absence of WAN links can lead to two-failure states, WAN link failure, or NSX Edge node failure, which can result in unrecoverable VMs and application failure at the remote site.

- Minimum bandwidth of 10 Mbps and latency of 100 ms is required between the central site and remote site. The network at the remote site must be able to reach the management network at the central site. DNS and NTP server must be available locally at or reachable from the remote site.
- See [VMware Configuration Maximums](#) for limitations related to VI workload domains at remote locations.
- See [VMware Cloud Foundation Edge Design Considerations](#) for more information about design options for deploying scalable edge solutions.

Change the VxRail Manager IP Address

In order to use the Workflow Optimization script to trigger VxRail APIs from the SDDC Manager VM, you must change the static IP address of the VxRail Manager to an IP address that is in the management network subnet.

- Ensure that a free IP address is available in the management network subnet
- Configure forward and reverse DNS settings for VxRail Manager
- The VxRail Manager static IP, 192.168.10.200, must be reachable and the UI available

1. Enter the following address in a web browser on your host `https://192.168.10.200/rest/vxm/api-doc.html#/operations/vl_network_vxm_post`.
2. Update the sample request body.

ip	Enter the new IP address for the VxRail Manager.
gateway	Enter the network gateway address for VxRail Manager.
netmask	Enter the subnet mask for VxRail Manager.
vlan_id	Enter the management network VLAN ID

3. Click **Send Request**.
4. Verify that the new IP address is reachable.

Update the VxRail Manager certificate. See [Update the VxRail Manager Certificate](#).

Update the VxRail Manager Certificate

After you change the VxRail Manager IP address to support using the Workflow Optimization script, you must update the VxRail Manager certificate.

Change the VxRail Manager IP Address

1. Using SSH, log in to VxRail Manager VM using the management IP address, with the user name `mystic` and default `mystic` password.
2. Type `su` to switch to the root account and enter the default root password.
3. Navigate to the `/mystic` directory.
4. Run the script:

```
./generate_ssl.sh VxRail-Manager-FQDN
```

Replace `VxRail-Manager-FQDN` with the VxRail Manager hostname.

Creating VxRail VI Workload Domains

You can create a VxRail VI workload domain using the SDDC Manager UI or using the Workflow Optimization script.

When you use the product UI, you complete the steps in the SDDC Manager UI. Starting with VMware Cloud Foundation 5.1, you can use the SDDC Manager UI to create a VI workload domain with advanced switch configurations.

Alternatively, you can use the Workflow Optimization script to create a VI workload domain. See [Create a VxRail VI Workload Domain Using the Workflow Optimization Script](#). The Workflow Optimization script supports using a JSON file for cluster configuration.

Create a VxRail VI Workload Domain in the SDDC Manager UI

Use the VxRail VI Configuration wizard to create a VI workload domain.

To create a VI workload domain that uses a static IP pool for the Host Overlay Network TEPs for L3 aware and stretch clusters, you must use the VMware Cloud Foundation API. See [Create a Domain](#) in the *VMware Cloud Foundation on Dell VxRail API Reference Guide*.

The SDDC Manager UI supports running multiple VxRail VI workload domain creation tasks in parallel.


1. In the SDDC Manager UI navigation, **Inventory > Workload Domains**.
2. Click **+ Workload Domain** and then select **VI-VxRail Virtual Infrastructure Setup**.
3. Make sure the prerequisites are met. See [Prerequisites for a Workload Domain](#). To continue, click **GET STARTED**.
4. Select the type of storage to use for this workload domain. Click **SELECT**.

Storage Selection

Select the type of storage you would like to use for this Workload Domain.

vSAN
Configure vSAN based workload domain.

Enable vSAN ESA

 vSAN ESA requires using vLCM images to manage clusters in workload domains.

VMFS on FC
Configure Fibre Channel based workload domain.

CANCEL

SELECT

NOTE

vSAN Express Storage Architecture (ESA) requires vSphere Lifecycle Manager images.

5. Provide the following information to complete the VxRail VI Configuration.

VxRail Manager	<ul style="list-style-type: none"> • VxRail Manager Hostname (must be an FQDN) CONNECT to VxRail Manager and confirm the SSL thumbprints of VxRail Manager.
----------------	---

Table continued on next page

Continued from previous page

	<ul style="list-style-type: none"> • VxRail Manager Admin Credentials <ul style="list-style-type: none"> – Admin Username – Admin Password – Confirm Admin Password • VxRail Manager Root Credentials <ul style="list-style-type: none"> – Root Username – Root Password – Confirm Root Password
General Info	<p>Provide basic information about the workload domain, including the SSO domain. When you create a VI workload domain, you can join it to the management domain's vCenter Single Sign-On domain or a new vCenter Single Sign-On domain that is not used by any other workload domain. Joining a new vCenter Single Sign-On domain enables a VI workload domain to be isolated from the other workload domains in your VMware Cloud Foundation instance. The vCenter Single Sign-On domain for a VI workload domain determines the local authentication space.</p> <ul style="list-style-type: none"> • Virtual Infrastructure Name - The name must be unique and contain between 3 and 20 characters. The VI name can include letters, numbers, and hyphens, but it cannot include spaces. • Datacenter Name • SSO domain <ul style="list-style-type: none"> – Create New SSO Domain <p style="text-align: center;">NOTE All components in the management domain must be upgraded to VMware Cloud Foundation 5.0 before you can create a new SSO domain.</p> <ul style="list-style-type: none"> – Join Management SSO Domain • Lifecycle Management - Select the Manage clusters in this workload domain using vLCM images check box to use vSphere Lifecycle Manager images. If you do not select the check box, the clusters in the workload domain use vSphere Lifecycle Manager baselines.

Table continued on next page

Continued from previous page

	<p>NOTE</p> <ul style="list-style-type: none"> vLCM images are managed by VxRail Manager. vSAN Express Storage Architecture (ESA) requires vSphere Lifecycle Manager images. Two-node clusters are not supported in a VI workload domain that uses vSphere Lifecycle Manager baselines. <p>If you are creating a new SSO domain, provide the following information:</p> <ul style="list-style-type: none"> Enter the domain name, for example <code>mydomain.local</code>. <p>NOTE</p> <p>Ensure that the domain name does not contain any upper-case letters.</p> <ul style="list-style-type: none"> Set the password for the SSO administrator account. This is the password for the user <code>administrator@your_domain_name</code>. Confirm the administrator password.
Host Selection	<p>Add ESXi hosts with similar or identical configurations across all cluster members, including similar or identical storage configurations. A minimum of 3 hosts are required.</p> <p>NOTE</p> <p>The Primary node is selected by default</p> <ol style="list-style-type: none"> Select the ESXi hosts to add and click Provide Host Details. Enter the FQDNs and passwords for the hosts. Click Resolve Hosts IP address. Click Next.
Cluster	<p>Enter a name for the first cluster in the new workload domain.</p> <p>The name must be unique and contain between 3 and 80 characters. The cluster name can include letters, numbers, and hyphens, and it can include spaces.</p>
Compute	<p>Provide information about the vCenter configuration.</p> <ul style="list-style-type: none"> vCenter FQDN (Must be a fully qualified domain name. (FQDN)) vCenter Subnet Mask vCenter Default Gateway vCenter Root Password Confirm vCenter Root Password

Table continued on next page

Continued from previous page

Networking	<p>Provide information about the NSX Manager cluster to use with the VI workload domain. If you already have an NSX Manager cluster for a different VI workload domain, you can reuse that NSX Manager cluster or create a new one.</p> <ul style="list-style-type: none"> • Create New NSX instance <p>NOTE</p> <ul style="list-style-type: none"> • You must create an NSX Manager instance if this is the first VI workload domain in your VMware Cloud Foundation instance. • You must create a new NSX Manager instance if your VI workload domain is joining a new SSO domain. <p>– Provide the NSX Manager cluster details:</p> <ul style="list-style-type: none"> • NSX Manager cluster FQDN • FQDNs for three NSX Managers nodes • Subnet mask • Default gateway • NSX Manager Admin password • Use Existing NSX instance <p>NOTE</p> <ul style="list-style-type: none"> • You cannot share an NSX Manager instance between VI workload domains that are in different SSO domains. • If you are creating a new SSO domain for the VI workload domain, the NSX Manager instance will be shared with one or more VI workload domains in different SSO domains. • In order to share an NSX Manager instance, the VI workload domains must use the same update method. The VI workload domains must both use vSphere Lifecycle Manager baselines or they must both use vSphere Lifecycle Manager images. <p>– Select the NSX Manager instance.</p> <p>NOTE NSX Managers for workload domains that are in the process of deploying are not able to be shared and do not appear in the list of available NSX Managers.</p>
------------	--

Table continued on next page

Continued from previous page

Switch Configuration	<p>Provide the distributed switch configuration to be applied to the hosts in the VxRail cluster. Select a predefined vSphere distributed switch (VDS) configuration profile or create a custom switch configuration.</p> <p>For custom switch configuration, specify:</p> <ul style="list-style-type: none"> • VDS name • MTU • Number of uplinks • Uplink to vmnic mapping <p>Click Configure Network Traffic to configure the following networks:</p> <ul style="list-style-type: none"> • Management • vMotion • vSAN • Host Discovery • System VM <p>For each network, specify:</p> <ul style="list-style-type: none"> • Distributed port group name • MTU • Load balancing policy • Active and standby links <p>For the NSX network, specify:</p> <ul style="list-style-type: none"> • Operational mode • Transport zone type • NSX-Overlay Transport Zone Name • For NSX Overlay, enter a VLAN ID and select the IP assignment type for the Host Overlay Network TEPs. <p>NOTE</p> <p>For DHCP, a DHCP server must be configured on the NSX host overlay (Host TEP) VLAN. When NSX creates TEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.</p> <p>For static IP Pool, you can re-use an existing IP pool or create a new one. Make sure the IP range includes enough IP addresses for the number of hosts that will use the static IP Pool. The number of IP addresses required depends on the number of pNICs on the ESXi hosts that are used for the vSphere Distributed Switch that handles host overlay networking. For example, a host with four pNICs that uses two pNICs for host overlay traffic requires two IP addresses in the static IP pool.</p>
----------------------	---

Table continued on next page

Continued from previous page

	<ul style="list-style-type: none"> • Teaming policy uplink mapping • NSX Uplink Profile Name • Teaming policy • Active and standby links <p>NOTE VDS configuration requires homogeneous host network adapters across all hosts. Only adapters of same enumeration across all hosts can be used for configuring VDS.</p>
Host Networks	<p>Configure the Host network details.</p> <ul style="list-style-type: none"> • Management Network: VLAN ID, CIDR, and Gateway • vSAN: VLAN ID, CIDR, Gateway, and IP Range • vMotion Network: VLAN ID, CIDR, Gateway, and IP Range • VM Management Network: Activate Same as Host Management or enter a VLAN ID, CIDR, and Gateway.
Licenses	<p>Select License Now or License Later.</p> <ul style="list-style-type: none"> • License Now: Select a license key for each of the components in the VI workload domain. • License Later: VMware Cloud Foundation components are deployed in evaluation mode. <p>IMPORTANT After your VI workload domain is created, you must switch to licensed mode by:</p> <ul style="list-style-type: none"> • Adding component license keys in the SDDC Manager UI. See Add a Component License Key in the SDDC Manager UI. Or, • Adding a solution license key in the vSphere Client. See Managing vSphere Licenses for information about using a solution license key for vCenter Server. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See Configure License Settings for a vSAN Cluster. <p>NOTE After you assign a solution key for vCenter Server, VMware NSX automatically uses that solution license key.</p>

Table continued on next page

Continued from previous page

Review	Review and confirm the Workload Domain settings.
Validation	Validates the configuration

6. On the **Validation** page, wait until all of the inputs have been successfully validated and then click **Finish**.
If validation is unsuccessful, you cannot proceed. Use the **Back** button to modify your settings and try again.
Create a VxRail VI Workload Domain task is triggered.

Create a VxRail VI Workload Domain Using the Workflow Optimization Script

You can create a VxRail VI workload domain using the Workflow Optimization script.

Make sure that the [Prerequisites for a Workload Domain](#) are met before using the Workflow Optimization script.

The Workflow Optimization script uses the VMware Cloud Foundation on Dell VxRail API to perform all of the steps to create a VI workload domain in one place. See [Create a Domain with Workflow Optimization](#) for more information about the API.

1. Download the .zip file for the Workflow Optimization script.
See <https://community.broadcom.com/vmware-code/viewdocument/vcf-on-vxrail-workflow-optimization-8>.
2. Unzip the file and copy the directory (`WorkflowOptimization-VCF-<version>`) to the `/home/vcf` directory on the SDDC Manager VM.
3. Remove the `storage` name/value from the JSON.
4. Using SSH, log in to the SDDC Manager VM as `vcf`.
5. In the `/home/vcf/WorkflowOptimization-VCF-<version>` directory, run `python vxrail_workflow_optimization_automator.py`.
6. Follow the prompts to create a VI workload domain.
The `README.md` file in the `WorkflowOptimization-VCF-<version>` directory provides detailed instructions on how to use the script.

Delete a VI Workload Domain

You can delete a VI workload domain from SDDC Manager UI.

- If remote vSAN datastores are mounted on a cluster in the VI workload domain, then the VI workload domain cannot be deleted. To delete such VI workload domains, you must first migrate any virtual machines from the remote datastore to the local datastore and then unmount the remote vSAN datastores from vCenter Server.
- If you require access after deleting a VI workload domain, back up the data. The datastores on the VI workload domain are destroyed when it is deleted.
- Migrate the virtual machines that you want to keep to another workload domain using cross vCenter vMotion.
- Delete any workload virtual machines created outside VMware Cloud Foundation before deleting the VI workload domain.
- Delete any NSX Edge clusters hosted on the VI workload domain. See [KB 78635](#).

Deleting a VI workload domain also removes the components associated with the VI workload domain from the management domain. This includes the vCenter Server instance and the NSX Manager cluster instances.

NOTE

If the NSX Manager cluster is shared with any other VI workload domains, it will not be deleted.

CAUTION

Deleting a workload domain is an irreversible operation. All clusters and virtual machines within the VI workload domain are deleted and the underlying datastores are destroyed.

It can take up to 20 minutes for a VI workload domain to be deleted. During this process, you cannot perform any operations on workload domains.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. Click the vertical ellipsis (three dots) next to the VI workload domain you want to delete and click **Delete Domain**.



Add Cluster
Add Edge Cluster
Delete Domain
Rename Domain

3. On the **Delete Workload Domain** dialog box, click **Delete Workload Domain**.

A message indicating that the VI workload domain is being deleted appears. When the removal process is complete, the VI workload domain is removed from the domains table.


If you delete an isolated VI workload domain that created an NSX Manager cluster that is shared with another isolated VI workload domain, you need to register NSX Manager as a relying partner to the remaining VI workload domain. See <https://kb.vmware.com/s/article/95445>.

View Workload Domain Details

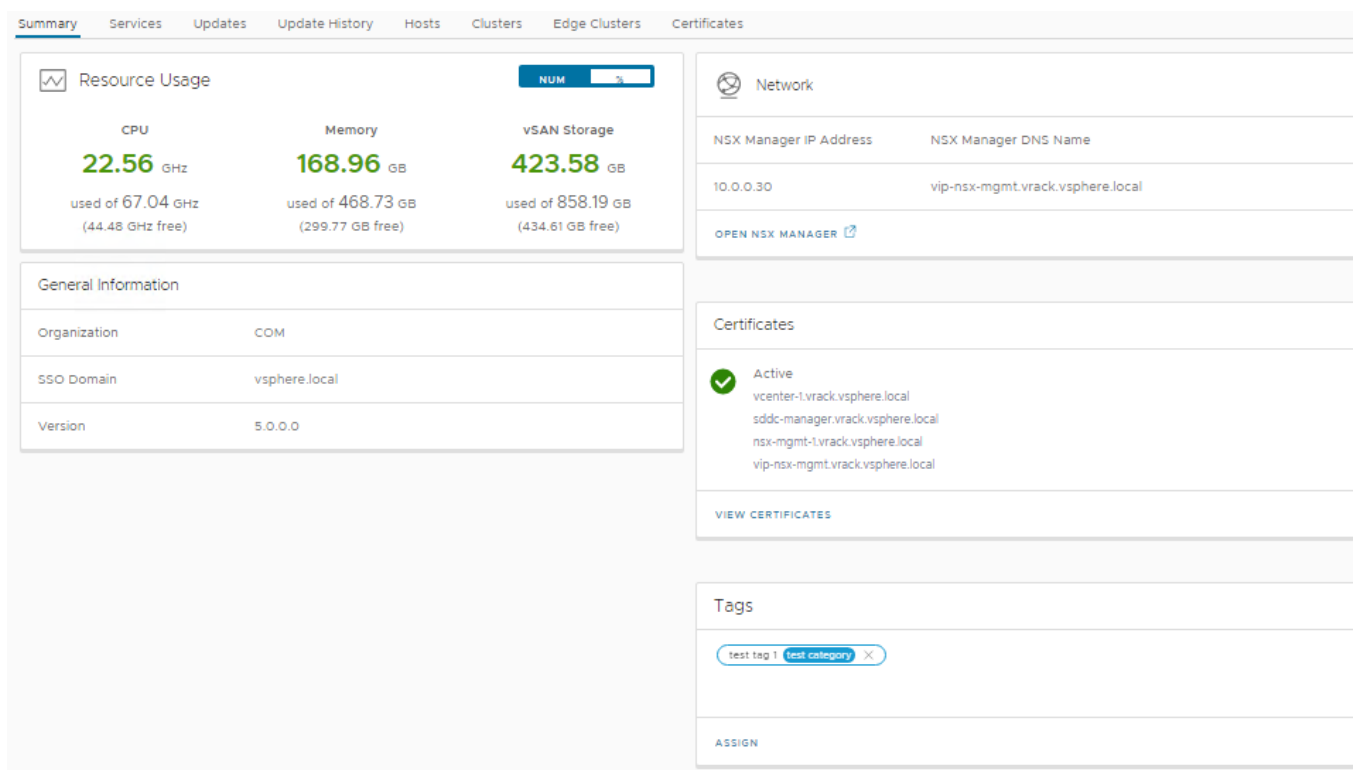
The Workload Domains page displays high level information about the workload domains in a VMware Cloud Foundation instance. CPU, memory, and storage utilized by the workload domain is also displayed here.

1. In the navigation pane, click **Inventory** › **Workload Domains**.

TIP

Click the show or hide columns icon  to view additional information about the workload domains, including the SSO domain.

2. In the workload domains table, click the name of the workload domain.



Tab	Information Displayed
Summary	Provides information about: <ul style="list-style-type: none"> Resource usage: CPU, memory, and storage resources for the workload domain. Network: NSX Manager IP address and DNS name. General information, including SSO domain. Certificates Tags
Services	SDDC software stack components deployed for the workload domain's virtual environment and their IP addresses. Click a component name to navigate to that aspect of the virtual environment. For example, click vCenter Server to reach the vSphere Client for that workload domain. All the capabilities of a VMware SDDC are available to you in the VI workload domain's environment, such as creating, provisioning, and deploying virtual machines, configuring the software-defined networking features, and so on.
Updates	Available updates for the workload domain.
Update History	Updates applied to this workload domain.
Hosts	Names, IP addresses, status, associated clusters, and capacity utilization of the hosts in the workload domain and the network pool they are associated with.
Clusters	Names of the clusters, number of hosts in the clusters, and their capacity utilization.
Edge Clusters	Names of the NSX Edge clusters, NSX Edge nodes, and their status.
Certificates	Default certificates for the VMware Cloud Foundation components. For more information, see Managing Certificates in VMware Cloud Foundation .

Expand a Workload Domain

You can expand a workload domain by adding a new VxRail cluster or adding hosts to an existing VxRail cluster.

To add a VxRail cluster to a workload domain, you can use the SDDC Manager UI or the Workflow Optimization script.

Method	Details
SDDC Manager UI	Supports most uses cases, including multiple vSphere distributed switches and advanced switch configuration.
Workflow Optimization script	Supports all SDDC Manager UI functionality. In addition, supports custom datastore names, remote vSAN datastores, and using a JSON file for cluster configuration. Download the script and refer to the <code>README.md</code> for instructions.

Add a VxRail Cluster to a Workload Domain Using the SDDC Manager UI

You can expand an existing workload domain by adding a VxRail cluster using the SDDC Manager UI.

- Image the workload domain nodes. For information on imaging the nodes, refer to Dell EMC VxRail documentation.
- The IP addresses and Fully Qualified Domain Names (FQDNs) for the ESXi hosts, VxRail Manager, and NSX Manager instances must be resolvable by DNS.
- If you are using DHCP for the NSX Host Overlay Network, a DHCP server must be configured on the NSX Host Overlay VLAN of the management domain. When VMware NSX creates TEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.
- [Change the VxRail Manager IP Address](#)
- [Update the VxRail Manager Certificate](#)

1. In the navigation pane, click **Inventory > Workload Domains**. The **Workload Domains** page displays information for all workload domains.
2. In the workload domains table, hover your mouse in the VxRail workload domain row. A set of three dots appears on the left of the workload domain name.
3. Click these three dots. Click **Add VxRail Cluster**.
4. Make sure the prerequisites are met. To continue, click **Get Started**.
5. Select the type of storage to use for this workload domain. Click **Select**.

For vSAN storage, you can enable vSAN ESA if the workload domain is using vSphere Lifecycle Manager images.

6. Provide the following information to Add VxRail Cluster to VI- VxRail.

VxRail Manager	<ul style="list-style-type: none"> • VxRail Manager Hostname (must be an FQDN) Click Connect and confirm the SSL fingerprint of the VxRail Manager. • VxRail Manager Admin Credentials <ul style="list-style-type: none"> – Admin Username – Admin Password – Confirm Admin Password • VxRail Manager Root Credentials <ul style="list-style-type: none"> – Root Username – Root Password – Confirm Root Password
----------------	---

Table continued on next page

Continued from previous page

Host Selection	<p>Add ESXi hosts with similar or identical configurations across all cluster members, including similar or identical storage configurations. A minimum of 3 hosts are required.</p> <p>NOTE The Primary node is selected by default</p> <ol style="list-style-type: none"> 1. Select the ESXi hosts to add and click Provide Host Details. 2. Enter the FQDNs and passwords for the hosts. 3. Click Resolve Hosts IP address. 4. Click Next.
Cluster	<p>Enter a name for the first cluster that will be created in this new workload domain. The name must be unique and contain between 3 and 80 characters. The cluster name can include letters, numbers, and hyphens, and it can include spaces.</p>
Switch Configuration	<p>Provide the distributed switch configuration to be applied to the hosts in the VxRail cluster. Select a predefined vSphere distributed switch (VDS) configuration profile or create a custom switch configuration.</p> <p>For custom switch configuration, specify:</p> <ul style="list-style-type: none"> • VDS name • MTU • Number of uplinks • Uplink to vmnic mapping <p>Click Configure Network Traffic to configure the following networks:</p> <ul style="list-style-type: none"> • Management • vMotion • vSAN • Host Discovery • System VM <p>For each network, specify:</p> <ul style="list-style-type: none"> • Distributed port group name • MTU • Load balancing policy • Active and standby links <p>For the NSX network, specify:</p> <ul style="list-style-type: none"> • Operational mode • Transport zone type • NSX-Overlay Transport Zone Name • For NSX Overlay, enter a VLAN ID and select the IP assignment type for the Host Overlay Network TEPs.

Table continued on next page

Continued from previous page

	<p>NOTE For DHCP, a DHCP server must be configured on the NSX host overlay (Host TEP) VLAN. When NSX creates TEPs for the VI workload domain, they are assigned IP addresses from the DHCP server.</p> <p>For static IP Pool, you can re-use an existing IP pool or create a new one. Make sure the IP range includes enough IP addresses for the number of hosts that will use the static IP Pool. The number of IP addresses required depends on the number of pNICs on the ESXi hosts that are used for the vSphere Distributed Switch that handles host overlay networking. For example, a host with four pNICs that uses two pNICs for host overlay traffic requires two IP addresses in the static IP pool.</p> <ul style="list-style-type: none"> • Teaming policy uplink mapping • NSX Uplink Profile Name • Teaming policy • Active and standby links <p>NOTE VDS configuration requires homogeneous host network adapters across all hosts. Only adapters of same enumeration across all hosts can be used for configuring VDS.</p>
Host Networks	<p>Configure the Host network details.</p> <ul style="list-style-type: none"> • Management Network: VLAN ID, CIDR, and Gateway • vSAN: VLAN ID, CIDR, Gateway, and IP Range • vMotion Network: VLAN ID, CIDR, Gateway, and IP Range • VM Network: Select Same as Host Management to use the host Management network information to create a new port group for VM Management traffic or enter a VLAN ID, CIDR, and Gateway for the new port group.
Licenses	<p>Select License Now or License Later.</p> <ul style="list-style-type: none"> • License Now: Select a license key for each of the components in the cluster. • License Later: VMware Cloud Foundation components are deployed in evaluation mode.

Table continued on next page

Continued from previous page

	<p>IMPORTANT</p> <p>After the cluster is created, you must switch to licensed mode by:</p> <ul style="list-style-type: none"> • Adding component license keys in the SDDC Manager UI. See Add a Component License Key in the SDDC Manager UI. Or, • Adding a solution license key in the vSphere Client. See Managing vSphere Licenses. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See Configure License Settings for a vSAN Cluster.
Review	Review and confirm the Workload Domain settings.
Validation	Validates the configuration

7. On the **Validation** page, wait until all of the inputs have been successfully validated.
If validation is unsuccessful, you cannot proceed. Use the **Back** button to modify your settings and try again.
8. Click **Finish**.
The add VxRail cluster task is triggered.

Add VxRail Hosts to a Cluster in VMware Cloud Foundation

You can add new hosts to an existing VxRail cluster to provide more capacity.

- Image the new node(s).
- Discover and add new node(s) to the cluster using the VxRail Manager plugin for vCenter Server. See the Dell documentation.

If the vSphere cluster hosts an NSX Edge cluster, you can only add new hosts with the same management, uplink, host TEP, and Edge TEP networks (L2 uniform) as the existing hosts.

If the cluster to which you are adding hosts uses a static IP pool for the Host Overlay Network TEPs, that pool must include enough IP addresses for the hosts you are adding. The number of IP addresses required depends on the number of pNICs on the ESXi hosts that are used for the vSphere Distributed Switch that handles host overlay networking. For example, a host with four pNICs that uses two pNICs for host overlay traffic requires two IP addresses in the static IP pool.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the workload domains table, click the name of the workload domain that you want to expand.
3. Click the **Clusters** tab.
4. Click the name of the cluster where you want to add a host.
5. Click **Actions** > **Add VxRail Hosts**.
6. Select the cluster expansion type.


This option only appears if the vSphere cluster hosts an NSX Edge cluster.

L2 Uniform	Select if all hosts you are adding to the vSphere cluster have the same management, uplink, host TEP, and Edge TEP networks as the existing hosts in the vSphere cluster.
------------	---

Table continued on next page

Continued from previous page

L2 non-uniform and L3	You cannot proceed if you any of the hosts you are adding to the vSphere cluster have different networks than the existing hosts in the vSphere cluster. VMware Cloud Foundation does not support adding hosts to L2 non-uniform and L3 vSphere clusters that host an NSX Edge cluster.
-----------------------	--

7. On the **Discovered Hosts** page, enter the SSH password for the host and click **Add**.
8. On the **Thumbprint Verification** page, click  to confirm the SSH thumbprints for the ESXi hosts.
9. On the **Validation** page, wait until all of the inputs have been successfully validated.
If validation is unsuccessful, you cannot proceed. Use the **Back** button to modify your settings and try again.
10. Click **Finish**.

Reduce a Workload Domain

You can reduce a workload domain by removing a host from a cluster in the workload domain or by deleting a cluster.

Remove a Host from a Cluster in a Workload Domain

You can remove a host from a cluster in a workload domain through the **Workload Domains** page in SDDC Manager UI.

Use the vSphere Client to make sure that there are no critical alarms on the cluster from which you want to remove the host.

When a host is removed, the vSAN members are reduced. Ensure that you have enough hosts remaining to facilitate the configured vSAN availability. Failure to do so might result in the datastore being marked as read-only or in data loss.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the workload domains table, click the name of the workload domain that you want to modify.
3. Click the **Clusters** tab.
4. Click the name of the cluster from which you want to remove a host.
5. Click the **Hosts** tab.
6. Select the host(s) to remove and click **Remove Selected Hosts**.
7. Click **Remove** to confirm the action.

The details page for the cluster appears with a message indicating that the host is being removed. When the removal process is complete, the host is removed from the hosts table and deleted from vCenter Server.

Delete a VxRail Cluster

You can delete a VxRail cluster from the management domain or from a VI workload domain. Datastores on the ESXi hosts in the deleted cluster are destroyed.

- If vSAN remote datastores are mounted on the cluster, the cluster cannot be deleted. To delete such clusters, you must first migrate any VMs from the remote datastore to the local datastore and then unmount the vSAN remote datastores from vCenter Server.
- Delete any workload VMs created outside of VMware Cloud Foundation before deleting the cluster.
- Migrate or backup the VMs and data on the datastore associated with the cluster to another location.
- Delete the NSX Edge clusters hosted on the VxRail cluster or shrink the NSX Edge cluster by deleting Edge nodes hosted on the VxRail cluster. You cannot delete Edge nodes if doing so would result in an Edge cluster with fewer than two Edge nodes. For information about deleting an NSX Edge cluster, see [KB 78635](#).

You cannot delete the last cluster in a workload domain. Instead, delete the workload domain.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
The Workload Domains page displays information for all workload domains.
2. Click the name of the workload domain that contains the cluster you want to delete.
3. Click the **Clusters** tab to view the clusters in the workload domain.
4. Hover your mouse in the cluster row you want to delete.
5. Click the three dots next to the cluster name and click **Delete VxRail Cluster**.
6. Click **Delete Cluster** to confirm that you want to delete the cluster.
The details page for the workload domain appears with a message indicating that the cluster is being deleted.
When the removal process is complete, the cluster is removed from the clusters table.

Rename a Workload Domain

You can rename any workload domain from within the SDDC Manager UI.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. Click the vertical ellipsis (three dots) in the Domain row for the workload domain you want to rename and click **Rename Domain**.
3. Enter a new name for the workload domain and click **Rename**.

vSphere Cluster Management

You can view vSphere cluster details from the SDDC Manager UI and rename the vSphere Cluster using the vSphere Client if required.

View vSphere Cluster Details

The cluster summary page displays high level information about the vSphere cluster as well as the hosts that form that cluster. CPU, memory, and storage utilization are also displayed.

1. In the navigation pane, click **Inventory** › **Workload Domain**.
2. In the workload domains table, click the name of a workload domain.
3. Click the **Clusters** tab.
4. In the clusters table, click the name of a vSphere cluster.

The screenshot shows the 'vi-cluster1' cluster detail page in the SDDC Manager UI. The page is in the 'Summary' tab and displays the following information:

- Resource Usage:** CPU usage is 8.94 GHz (used of 50.28 GHz, 41.34 GHz free). Memory usage is 56.12 GB (used of 351.55 GB, 295.43 GB free).
- Storage:** Local vSAN storage is used, showing 36.43 GB used of 858.23 GB (821.8 GB free).
- Datastore Name:** vi-cluster1-vSanDatastore
- FTT:** 0
- Tags:** No tags assigned.

The cluster detail page appears. The tabs on the page display additional information as described in the table below.

Tab	Information Displayed
Summary	Displays information about resource usage, storage, and cluster tags.
Hosts	Details about the ESXi hosts in the vSphere cluster. You can click a name in the FQDN column to access the host summary page.

You can add or remove a host, or access the vSphere Client from this page.

Rename a Cluster in the SDDC Manager UI

You can rename a cluster managed by SDDC Manager in a Management Workload Domain. The SDDC Manager UI is updated with the new name.

Ensure that you do not rename a cluster in the following conditions:

- When the cluster belongs to a workflow that is in progress.
- When the cluster belongs to a failed VI workload domain workflow, cluster workflow or host workflow. If you try to rename a cluster that belongs to a failed workflow, restart of the failed workflow will not be supported.

1. On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
2. Click a workload domain.
3. Under the **Clusters** tab, click a cluster that you want to rename.
4. On the right side of the cluster's name, click **ACTIONS > Rename Cluster**.

You can also click the vertical ellipsis (three dots) in the clusters table for the cluster you want to rename and click **Rename Cluster**.

The **Rename Cluster** window appears.

- In the **New Cluster Name** textbox, enter a new name for the cluster and click **RENAME**.
- Click **DONE**.

In the **Tasks** panel, you can see the description and track the status of your newly renamed cluster.

Tag Management

A tag is a label that you can apply to objects in the vSphere inventory. You can use tags to capture a variety of metadata about your vSphere inventory and to organize and retrieve objects quickly. You create tags and categories in the vSphere Client and then assign or remove tags for your workload domains, clusters, and hosts in the SDDC Manager UI.

See [vSphere Tags](#) for more information about how to create and manage tags and categories.

If multiple vCenter Server instances in your VMware Cloud Foundation deployment are configured to use Enhanced Linked Mode, tags and tag categories are replicated across all these vCenter Server instances. This is the case for all VI workload domains that are joined to the same SSO domain as the management domain. Isolated VI workload domains, that do not share the management SSO domain, do not share its tags and categories.

Tag a Workload Domain

You can assign a tag to your Workload Domain from the SDDC Manager UI by performing the following steps:

- On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** and click the Workload Domain.
- Under the **Summary** › **Tags** tile window, click **ASSIGN**.

Assign Tag

[OPEN VSPHERE TAG MANAGEMENT](#) ✕

Assign tags to the workload domain "nsxt-vi". To create new tags or new tag categories, go to vSphere Tag Management.

<input type="checkbox"/>	Tag name	↑ ▼	Tag category	▼	Tag associations per category	▼
<input type="checkbox"/>	workload-domain-production		Workload Domains		One	
<input type="checkbox"/>	workload-domain-test		Workload Domains		One	
						1 - 2 of 2 Tags

CANCEL

ASSIGN

- Select tags and click **ASSIGN**.

NOTE

If there are no tags shown in the **Assign Tag** window, click **OPEN VSPHERE TAG MANAGEMENT** that redirects you to vSphere Client, to create new tags and tag categories. See [vSphere Tags](#) for more information on the tagging functionality.

Remove a Tag from your Workload Domain

You can remove a tag from your workload domain in the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** and click the workload domain.
2. Under the **Summary** › **Tags** tile window, you will see tags listed with a cross mark beside the tag names.
3. Click the cross mark of a tag that you want to remove in the **Tags** tile window.
4. The **Remove Tag** window appears. Click **REMOVE**.

Tag a Cluster

You can assign a tag to your cluster from the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** › **Workload Domain** › **Clusters** tab and click on the cluster.
2. Under the **Summary** › **Tags** tile window, click **ASSIGN**.

Assign Tag

[OPEN VSPHERE TAG MANAGEMENT](#) ✕

Assign tags to the cluster "vi-cluster1". To create new tags or new tag categories, go to vSphere Tag Management.

<input type="checkbox"/>	Tag name	Tag category	Tag associations per category
<input type="checkbox"/>	dev-cluster	Cluster Users	One
<input type="checkbox"/>	marketing-cluster	Cluster Users	One

1 - 2 of 2 Tags

CANCEL

ASSIGN

3. Select tags and click **ASSIGN**.

NOTE

If there are no tags shown in the **Assign Tag** window, click **OPEN VSPHERE TAG MANAGEMENT** that redirects you to vSphere Client, to create new tags and tag categories. See [vSphere Tags](#) for more information on the tagging functionality.

Remove a Tag from your Cluster

You can remove a tag from your cluster in the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Workload Domains** › **Management** › **Workload Domain** › **Clusters** tab and click on the cluster.
2. Under the **Summary** › **Tags** tile window, you will see tags listed with a cross mark beside the tag names.
3. Click the cross mark of a tag that you want to remove in the **Tags** tile window.
4. The **Remove Tag** window appears. Click **REMOVE**.

Tag a Host

You can assign a tag to your host from the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Hosts** tab and click on the host.
2. Under the **Summary** › **Tags** tile window, click **ASSIGN**.

Assign Tag

[OPEN VSPHERE TAG MANAGEMENT](#) ✕

Assign tags to the host "esxi-1.vrack.vsphere.local". To create new tags or new tag categories, go to vSphere Tag Management.

<input type="checkbox"/>	Tag name	↑ ▼	Tag category	▼	Tag associations per category	▼
<input type="checkbox"/>	mgmt-hosts		Workload Domain Hosts		One	
<input type="checkbox"/>	vi-wld-hosts		Workload Domain Hosts		One	
						1 - 2 of 2 Tags

CANCEL

ASSIGN

3. Select tags and click **ASSIGN**.

NOTE

If there are no tags shown in the **Assign Tag** window, click **OPEN VSPHERE TAG MANAGEMENT** that redirects you to vSphere Client, to create new tags and tag categories. See [vSphere Tags](#) for more information on the tagging functionality.

Remove a Tag from your Host

You can remove a tag from your host in the SDDC Manager UI by performing the following steps:

1. On the SDDC Manager UI, click **Inventory** › **Hosts** tab and click on the host.
2. Under the **Summary** › **Tags** tile window, you will see tags listed with a cross mark beside the tag names.
3. Click the cross mark of a tag that you want to remove in the **Tags** tile window.
4. The **Remove Tag** window appears. Click **REMOVE**.

Managing NSX Edge Clusters in VMware Cloud Foundation

An NSX Edge cluster with 2-tier routing provides north-south routing and network services in the management domain and VI workload domains. Add multiple NSX Edge clusters to a workload domain for scalability and resiliency.

An NSX Edge cluster is a logical grouping of NSX Edge nodes run on a vSphere cluster. NSX supports a 2-tier routing model.

Component	Connectivity	Description
Tier-0 logical router	Northbound	The tier-0 logical router connects to one or more physical routers or layer 3 switches and serves as a gateway to the physical infrastructure.
	Southbound	The tier-0 logical router connects to one or more tier-1 logical routers or

Table continued on next page

Continued from previous page

Component	Connectivity	Description
		directly to one or more logical switches.
Tier-1 logical router	Northbound	The tier-1 logical router connects to a tier-0 logical router.
	Southbound	The tier-1 logical router connects to one or more logical switches.

By default, workload domains do not include any NSX Edge clusters and workloads are isolated, unless VLAN-backed networks are configured in vCenter Server. Add one or more NSX Edge clusters to a workload domain to provide software-defined routing and network services.

NOTE

You must create an NSX Edge cluster on the default management vSphere cluster in order to deploy VMware Aria Suite products.

You can add multiple NSX Edge clusters to the management or the VI workload domains for scalability and resiliency. For VMware Cloud Foundation configuration maximums refer to the [VMware Configuration Maximums](#) website.

NOTE

Unless explicitly stated in this matrix, VMware Cloud Foundation supports the configuration maximums of the underlying products. Refer to the individual product configuration maximums as appropriate.

The north-south routing and network services provided by an NSX Edge cluster created for a workload domain are shared with all other workload domains that use the same NSX Manager cluster.

Prerequisites for an NSX Edge Cluster

Before you deploy an NSX Edge cluster on a workload domain, review the prerequisites.

- The workload domain must have NSX deployed.
- Verify that separate VLANs and subnets are available for the NSX host overlay VLAN and NSX Edge overlay VLAN. You cannot use DHCP for the NSX Edge overlay VLAN.
- Verify that the NSX host overlay VLAN and NSX Edge overlay VLAN are routed to each other.
- For dynamic routing, set up two Border Gateway Protocol (BGP) peers on Top of Rack (ToR) switches with an interface IP, BGP autonomous system number (ASN), and BGP password.
- Reserve a BGP ASN to use for the NSX Edge cluster's Tier-0 gateway.
- Verify that DNS entries for the NSX Edge nodes are populated in the customer-managed DNS server.
- The vSphere cluster hosting an NSX Edge cluster must include hosts with identical management, uplink, NSX Edge overlay TEP, and NSX Edge overlay TEP networks (L2 uniform).
- The management network and management network gateway for the NSX Edge nodes must be reachable from the NSX host overlay and NSX Edge overlay VLANs.

NOTE

VMware Cloud Foundation 4.5 and later support deploying an NSX Edge cluster on a vSphere cluster that is stretched. Edge nodes are placed on ESXi hosts in the first availability zone (AZ1) during NSX Edge cluster deployment.

Deploy an NSX Edge Cluster

Deploy an NSX Edge cluster to provide north-south routing and network services to a workload domain.

See [Prerequisites for an NSX Edge Cluster](#).

SDDC Manager does not enforce rack failure resiliency for NSX Edge clusters. Make sure that the number of NSX Edge nodes that you add to an NSX Edge cluster, and the vSphere clusters to which you deploy the NSX Edge nodes, are sufficient to provide NSX Edge routing services in case of rack failure.

After you create an NSX Edge cluster, you can use SDDC Manager to expand or shrink it by adding or deleting NSX Edge nodes.

NOTE

If you deploy the NSX Edge cluster with the incorrect settings or need to delete an NSX Edge cluster for another reason, see [KB 78635](#).

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Select **Actions** > **Add Edge Cluster**.
4. Verify the prerequisites, select **Select All**, and click **Begin**.
5. Enter the configuration settings for the NSX Edge cluster and click **Next**.

Setting	Description
Edge Cluster Name	Enter a name for the NSX Edge cluster.
MTU	Enter the MTU for the NSX Edge cluster. The MTU can be 1600-9000.
Tier-0 Router Name	Enter a name for the tier-0 gateway.
Tier-1 Router Name	Enter a name for the tier-1 gateway.
Edge Cluster Profile Type	Select Default or, if your environment requires specific Bidirectional Forwarding Detection (BFD) configuration, select Custom .
Edge Cluster Profile Name	Enter an NSX Edge cluster profile name. (Custom Edge cluster profile only)
BFD Allowed Hop	Enter the number of multi-hop Bidirectional Forwarding Detection (BFD) sessions allowed for the profile. (Custom Edge cluster profile only)
BFD Declare Dead Multiple	Enter the number of number of times the BFD packet is not received before the session is flagged as down. (Custom Edge cluster profile only)
BFD Probe Interval (milliseconds)	BFD is detection protocol used to identify the forwarding path failures. Enter a number to set the interval timing for BFD to detect a forwarding path failure. (Custom Edge cluster profile only)
Standby Relocation Threshold (minutes)	Enter a standby relocation threshold in minutes. (Custom Edge cluster profile only)
Edge Root Password	Enter and confirm the password to be assigned to the root account of the NSX Edge appliance.
Edge Admin Password	Enter and confirm the password to be assigned to the admin account of the NSX Edge appliance.
Edge Audit Password	Enter and confirm the password to be assigned to the audit account of the NSX Edge appliance.

NSX Edge cluster passwords must meet the following requirements:

- At least 12 characters
- At least one lower-case letter
- At least one upper-case letter
- At least one digit

- At least one special character (!, @, ^, =, *, +)
- At least five different characters
- No dictionary words
- No palindromes
- More than four monotonic character sequence is not allowed

6. Specify the use case details and click **Next**.

Setting	Description
Use Case	<ul style="list-style-type: none"> • Select Kubernetes - Workload Management to create an NSX Edge cluster that complies with the requirements for deploying vSphere IaaS Control Plane. See VMware Cloud Foundation with VMware Tanzu . If you select this option, you cannot modify the NSX Edge form factor or Tier-0 service high availability settings. • Select Application Virtual Networks to create an NSX Edge cluster that complies with the requirements deploying VMware Aria Suite components. See Deploying Application Virtual Networks in VMware Cloud Foundation. <p style="text-align: center;">NOTE Management domain only.</p> <ul style="list-style-type: none"> • Select Custom if you want an NSX Edge cluster with a specific form factor or Tier-0 service high availability setting.
Edge Form Factor	<ul style="list-style-type: none"> • Small: 4 GB memory, 2 vCPU, 200 GB disk space. The NSX Edge Small VM appliance size is suitable for lab and proof-of-concept deployments. • Medium: 8 GB memory, 4 vCPU, 200 GB disk space. The NSX Edge Medium appliance size is suitable for production environments with load balancing. • Large: 32 GB memory, 8 vCPU, 200 GB disk space. The NSX Edge Large appliance size is suitable for production environments with load balancing. • XLarge: 64 GB memory, 16 vCPU, 200 GB disk space. The NSX Edge Extra Large appliance size is suitable for production environments with load balancing.
Tier-0 Service High Availability	<p>In the active-active mode, traffic is load balanced across all members. In active-standby mode, all traffic is processed by an elected active member. If the active member fails, another member is elected to be active. Workload Management requires Active-Active.</p> <p>Some services are only supported in Active-Standby: NAT, load balancing, stateful firewall, and VPN. If you select Active-Standby, use exactly two NSX Edge nodes in the NSX Edge cluster.</p>

Table continued on next page

Continued from previous page

Setting	Description
Tier-0 Routing Type	Select Static or EBGP to determine the route distribution mechanism for the tier-0 gateway. If you select Static , you must manually configure the required static routes in NSX Manager. If you select EBGP , VMware Cloud Foundation configures eBGP settings to allow dynamic route distribution.
ASN	Enter an autonomous system number (ASN) for the NSX Edge cluster. (for EBGP only)

7. Enter the configuration settings for the first NSX Edge node and click **Add Edge Node**.

Setting	Description
Edge Node Name (FQDN)	Enter the FQDN for the NSX Edge node. Each node must have a unique FQDN.
Cluster	<p>Select a vSphere cluster to host the NSX Edge node. You can select a standard vSphere cluster or a stretched vSphere cluster, but all the NSX Edge nodes in an NSX Edge cluster must be hosted on vSphere clusters of the same type.</p> <p>NOTE If the vSphere cluster you select already hosts management virtual machines that are connected to the host Management port group, the VM Management Portgroup VLAN and VM Management Portgroup VLAN settings are not available.</p>
Cluster Type	<p>Select L2 Uniform if all hosts in the vSphere cluster have identical management, uplink, host TEP, and Edge TEP networks.</p> <p>Select L2 non-uniform and L3 if any of the hosts in the vSphere cluster have different networks.</p> <p>IMPORTANT VMware Cloud Foundation does not support Edge cluster creation on L2 non-uniform and L3 vSphere clusters.</p>
First NSX VDS Uplink	<p>Click Advanced Cluster Settings to map the first NSX Edge node uplink network interface to a physical NIC on the host, by specifying the ESXi uplink. The default is <code>uplink1</code>.</p> <p>When you create an NSX Edge cluster, SDDC Manager creates two trunked VLAN port groups. The information you enter here determines the active uplink on the first VLAN port group. If you enter <code>uplink3</code>, then <code>uplink3</code> is the active uplink and the uplink you specify for the second NSX VDS uplink is the standby uplink.</p> <p>The uplink must be prepared for overlay use.</p>

Table continued on next page

Continued from previous page

Setting	Description
Second NSX VDS Uplink	<p>Click Advanced Cluster Settings to map the second NSX Edge node uplink network interface to a physical NIC on the host, by specifying the ESXi uplink. The default is <code>uplink2</code>.</p> <p>When you create an NSX Edge cluster, SDDC Manager creates two trunked VLAN port groups. The information you enter here determines the active uplink on the second VLAN port group. If you enter <code>uplink4</code>, then <code>uplink4</code> is the active uplink and the uplink you specify for the first NSX VDS uplink is the standby uplink.</p> <p>The uplink must be prepared for overlay use.</p>
Management IP (CIDR)	Enter the management IP for the NSX Edge node in CIDR format. Each node must have a unique management IP.
Management Gateway	Enter the IP address for the management network gateway.
VM Management Portgroup VLAN	<p>If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group VLAN is displayed and cannot be edited.</p> <p>If the VM Management port group does not exist on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, enter a VLAN ID to create a new VM Management port group or click Use ESXi Management VMK's VLAN to use the host Management Network VLAN to create a new VM Management port group.</p>
VM Management Portgroup Name	<p>If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group name is displayed and cannot be edited.</p> <p>Otherwise, type a name for the new port group.</p>
Edge TEP 1 IP (CIDR)	<p>Enter the CIDR for the first NSX Edge TEP. Each node must have a unique Edge TEP 1 IP.</p> <p>NOTE It is possible to configure Edge TEPs using an NSX IP pool instead of static addresses. IP pools may only be specified when using the VCF API only, not the UI.</p>
Edge TEP 2 IP (CIDR)	Enter the CIDR for the second NSX Edge TEP. Each node must have a unique Edge TEP 2 IP. The Edge TEP 2 IP must be different than the Edge TEP 1 IP.
Edge TEP Gateway	Enter the IP address for the NSX Edge TEP gateway.
Edge TEP VLAN	Enter the NSX Edge TEP VLAN ID.
First Tier-0 Uplink VLAN	Enter the VLAN ID for the first uplink.

Table continued on next page

Continued from previous page

Setting	Description
	This is a link from the NSX Edge node to the first uplink network.
First Tier-0 Uplink Interface IP (CIDR)	Enter the CIDR for the first uplink. Each node must have unique uplink interface IPs.
Peer IP (CIDR)	Enter the CIDR for the first uplink peer. (EBGP only)
Peer ASN	Enter the ASN for the first uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only).
Second Tier-0 Uplink VLAN	Enter the VLAN ID for the second uplink. This is a link from the NSX Edge node to the second uplink network.
Second Tier-0 Uplink Interface IP (CIDR)	Enter the CIDR for the second uplink. Each node must have unique uplink interface IPs. The second uplink interface IP must be different than the first uplink interface IP.
Peer IP (CIDR)	Enter the CIDR for the second uplink peer. (EBGP only)
ASN Peer	Enter the ASN for the second uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only).

8. Click **Add More Edge Nodes** to enter configuration settings for additional NSX Edge nodes.

A minimum of two NSX Edge nodes is required. NSX Edge cluster creation allows up to 8 NSX Edge nodes if the Tier-0 Service High Availability is Active-Active and two NSX Edge nodes per NSX Edge cluster if the Tier-0 Service High Availability is Active-Standby.

NOTE

All Edge nodes in the NSX Edge cluster must use the same VM Management port group VLAN and name.

9. When you are done adding NSX Edge nodes, click **Next**.
10. Review the summary and click **Next**.
SDDC Manager validates the NSX Edge node configuration details.
11. If validation fails, use the **Back** button to edit your settings and try again.

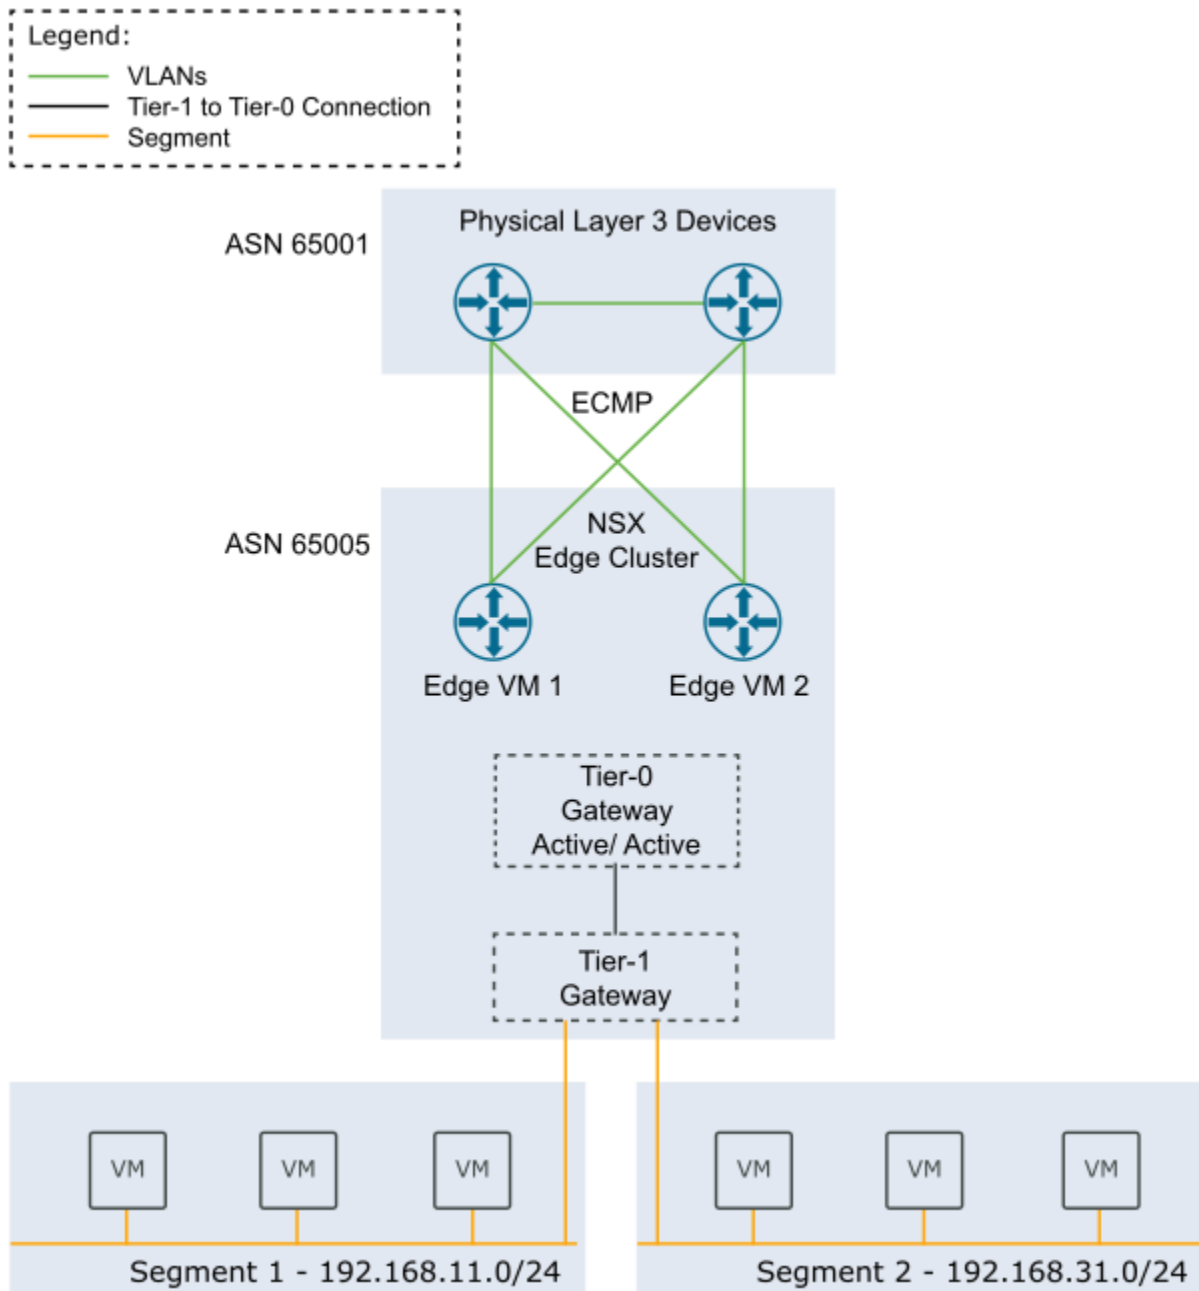
To edit or delete any of the NSX Edge nodes, click the three vertical dots next to an NSX Edge node in the table and select an option from the menu.

12. If validation succeeds, click **Finish** to create the NSX Edge cluster.

You can monitor progress in the Tasks panel.

The following example shows a scenario with sample data. You can use the example to guide you in creating NSX Edge clusters in your environment. Refer to the [Planning and Preparation Workbook](#) for a complete list of sample values for creating an NSX Edge cluster.

Figure 23: Two-node NSX Edge cluster in a single rack



In NSX Manager, you can create segments connected to the NSX Edge cluster's tier-1 gateway. You can connect workload virtual machines to these segments to provide north-south and east-west connectivity.

Add Edge Nodes to an NSX Edge Cluster

You can add NSX Edge nodes to an NSX Edge Cluster that you created with SDDC Manager.

- Verify that separate VLANs and subnets are available for the NSX host overlay VLAN and NSX Edge overlay VLAN. You cannot use DHCP for the NSX Edge overlay VLAN.
- Verify that the NSX host overlay VLAN and NSX Edge overlay VLAN are routed to each other.

- For dynamic routing, set up two Border Gateway Protocol (BGP) peers on Top of Rack (ToR) switches with an interface IP, BGP autonomous system number (ASN), and BGP password.
- Reserve a BGP ASN to use for the NSX Edge cluster's Tier-0 gateway.
- Verify that DNS entries for the NSX Edge nodes are populated in the customer-managed DNS server.
- The vSphere cluster hosting the NSX Edge nodes must include hosts with identical management, uplink, NSX Edge overlay TEP, and NSX Edge overlay TEP networks (L2 uniform).
- The vSphere cluster hosting the NSX Edge nodes must have the same pNIC speed for NSX-enabled VDS uplinks chosen for Edge overlay.
- All NSX Edge nodes in an NSX Edge cluster must use the same set of NSX-enabled VDS uplinks. These uplinks must be prepared for overlay use.
- The NSX Edge cluster must be **Active**.
- The NSX Edge cluster must be hosted on one or more vSphere clusters from the same workload domain.

You might want to add NSX Edge nodes to an NSX Edge cluster, for:

- Rack failure resiliency
- When the Tier-0 Service High Availability is Active-Standby and you require more than two NSX Edge nodes for services.

NOTE

Only two of the NSX Edge nodes can have uplink interfaces, but you can add more nodes without uplink interfaces.

- When the Tier-0 Service High Availability is Active-Active and you require more than 8 NSX Edge nodes for services.
- When you add Supervisor Clusters to a Workload Management workload domain and need to support additional tier-1 gateways and services.

The available configuration settings for a new NSX Edge node vary based on:

- The Tier-0 Service High Availability setting (Active-Active or Active-Standby) of the NSX Edge cluster.
- The Tier-0 Routing Type setting (static or EBGp) of the NSX Edge cluster.
- Whether the new NSX Edge node is going to be hosted on the same vSphere cluster as the existing NSX Edge nodes (in-cluster) or on a different vSphere cluster (cross-cluster).

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Click the **Edge Clusters** tab.
4. Click the vertical ellipsis menu for the Edge Cluster you want to expand and select **Expand Edge Cluster**.
5. Verify the prerequisites, select **Select All**, and click **Begin**.
6. Enter and confirm the passwords for the NSX Edge cluster.
7. Enter a name to create a new tier-1 gateway.
8. Enter the configuration settings for the new NSX Edge node and click **Add Edge Node**.

Setting	Description
Edge Node Name (FQDN)	Enter the FQDN for the NSX Edge node. Each node must have a unique FQDN.
Cluster	Select a vSphere cluster to host the NSX Edge node. If the workload domain has multiple vSphere clusters, you can select the vSphere cluster hosting the existing NSX Edge nodes (in-cluster expansion) or select a different vSphere cluster to host the new NSX Edge nodes (cross-cluster expansion).

Table continued on next page

Continued from previous page

Setting	Description
	<p>NOTE If the vSphere cluster you select already hosts management virtual machines that are connected to the host Management port group, the VM Management Portgroup VLAN and VM Management Portgroup VLAN settings are not available.</p>
Cluster Type	<p>Select L2 Uniform if all hosts in the vSphere cluster have identical management, uplink, host TEP, and Edge TEP networks. Select L2 non-uniform and L3 if any of the hosts in the vSphere cluster have different networks.</p> <p>IMPORTANT VMware Cloud Foundation does not support Edge cluster creation on L2 non-uniform and L3 vSphere clusters.</p>
Management IP (CIDR)	Enter the management IP for the NSX Edge node in CIDR format. Each node must have a unique management IP.
Management Gateway	Enter the IP address for the management network gateway.
VM Management Portgroup VLAN	<p>For in-cluster expansion, the new Edge node uses the same VM Management port group VLAN as the other Edge nodes in the Edge cluster.</p> <p>For cross-cluster expansion:</p> <ul style="list-style-type: none"> • If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group VLAN is displayed and cannot be edited. • If the VM Management port group does not exist on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, enter a VLAN ID to create a new VM Management port group or click Use ESXi Management VMK's VLAN to use the host Management Network VLAN for the VM Management port group.
VM Management Portgroup Name	<p>For in-cluster expansion, the new Edge node uses the same VM Management port group name as the other Edge nodes in the Edge cluster.</p> <p>For cross-cluster expansion:</p> <ul style="list-style-type: none"> • If the VM Management port group exists on the vSphere distributed switch of the vSphere cluster that you selected to host the Edge node, then the VM Management port group name is displayed and cannot be edited. • Otherwise, type a name for the port group.

Table continued on next page

Continued from previous page

Setting	Description
Edge TEP 1 IP (CIDR)	Enter the CIDR for the first NSX Edge TEP. Each node must have a unique Edge TEP 1 IP.
Edge TEP 2 IP (CIDR)	Enter the CIDR for the second NSX Edge TEP. Each node must have a unique Edge TEP 2 IP. The Edge TEP 2 IP must be different than the Edge TEP 1 IP.
Edge TEP Gateway	Enter the IP address for the NSX Edge TEP gateway.
Edge TEP VLAN	Enter the NSX Edge TEP VLAN ID.
First NSX VDS Uplink	<p>Specify an ESXi uplink to map the first NSX Edge node uplink network interface to a physical NIC on the host. The default is <code>uplink1</code>.</p> <p>The information you enter here determines the active uplink on the first VLAN port group used by the NSX Edge node. If you enter <code>uplink3</code>, then <code>uplink3</code> is the active uplink and the uplink you specify for the second NSX VDS uplink is the standby uplink.</p> <p>(cross-cluster only)</p> <p>NOTE For in-cluster NSX Edge cluster expansion, new NSX Edge nodes use the same NSX VDS uplinks as the other Edge nodes hosted on the vSphere cluster.</p>
Second NSX VDS Uplink	<p>Specify an ESXi uplink to map the second NSX Edge node uplink network interface to a physical NIC on the host. The default is <code>uplink2</code>.</p> <p>The information you enter here determines the active uplink on the second VLAN port group used by the NSX Edge node. If you enter <code>uplink4</code>, then <code>uplink4</code> is the active uplink and the uplink you specify for the first NSX VDS uplink is the standby uplink.</p> <p>(cross-cluster only)</p> <p>NOTE For in-cluster NSX Edge cluster expansion, new NSX Edge nodes use the same NSX VDS uplinks as the other Edge nodes hosted on the vSphere cluster.</p>
Add Tier-0 Uplinks	Optional. Click Add Tier-0 Uplinks to add tier-0 uplinks. (Active-Active only)
First Tier-0 Uplink VLAN	<p>Enter the VLAN ID for the first uplink. This is a link from the NSX Edge node to the first uplink network.</p> <p>(Active-Active only)</p>
First Tier-0 Uplink Interface IP (CIDR)	Enter the CIDR for the first uplink. Each node must have unique uplink interface IPs.

Table continued on next page

Continued from previous page

Setting	Description
	(Active-Active only)
Peer IP (CIDR)	Enter the CIDR for the first uplink peer. (EBGP only)
Peer ASN	Enter the ASN for the first uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only)
Second Tier-0 Uplink VLAN	Enter the VLAN ID for the second uplink. This is a link from the NSX Edge node to the second uplink network. (Active-Active only)
Second Tier-0 Uplink Interface IP(CIDR)	Enter the CIDR for the second uplink. Each node must have unique uplink interface IPs. The second uplink interface IP must be different than the first uplink interface IP. (Active-Active only)
Peer IP (CIDR)	Enter the CIDR for the second uplink peer. (EBGP only)
ASN Peer	Enter the ASN for the second uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only)

9. Click **Add More Edge Nodes** to enter configuration settings for additional NSX Edge nodes.

An NSX Edge cluster can contain a maximum of 10 NSX Edge nodes.

- For an NSX Edge cluster with a Tier-0 Service High Availability setting of Active-Active, up to 8 of the NSX Edge nodes can have uplink interfaces.
- For an NSX Edge cluster with a Tier-0 Service High Availability setting of Active-Standby, up to 2 of the NSX Edge nodes can have uplink interfaces.

10. When you are done adding NSX Edge nodes, click **Next**.

11. Review the summary and click **Next**.

SDDC Manager validates the NSX Edge node configuration details.

12. If validation fails, use the **Back** button to edit your settings and try again.

To edit or delete any of the NSX Edge nodes, click the three vertical dots next to an NSX Edge node in the table and select an option from the menu.

13. If validation succeeds, click **Finish** to add the NSX Edge node(s) to the NSX Edge cluster.

You can monitor progress in the Tasks panel.

Remove Edge Nodes from an NSX Edge Cluster

You can remove NSX Edge nodes from an NSX Edge Cluster that you created with SDDC Manager if you need to scale down to meet business needs.

- The NSX Edge cluster must be available in the SDDC Manager inventory and must be **Active**.
- The NSX Edge node must be available in the SDDC Manager inventory.
- The NSX Edge cluster must be hosted on one or more vSphere clusters from the same workload domain.

- The NSX Edge cluster must contain more than two NSX Edge nodes.
- The NSX Edge cluster must not be federated or stretched.
- If the NSX Edge cluster was deployed with a Tier-0 Service High Availability of Active-Active, the NSX Edge cluster must contain two or more NSX Edge nodes with two or more Tier-0 routers (SR component) after the NSX Edge nodes are removed.
- If selected edge cluster was deployed with a Tier-0 Service High Availability of Active-Standby, you cannot remove NSX Edge nodes that are the active or standby node for the Tier-0 router.

For information about deleting an NSX Edge cluster, see [KB 78635](#).

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Click the **Edge Clusters** tab.
4. Click the vertical ellipsis menu for the Edge Cluster you want to expand and select **Shrink Edge Cluster**.
5. Select the Edge node(s) to remove and click **Next**.
6. Review the summary and click **Next**.
SDDC Manager validates the request.
7. If validation fails, use the **Back** button to edit your settings and try again.

NOTE

You cannot remove the active and standby Edge nodes of a Tier-1 router at the same time. You can remove one and then remove the other after the first operation is complete.

8. If validation succeeds, click **Finish** to remove the NSX Edge node(s) from the NSX Edge cluster.

You can monitor progress in the Tasks panel.

Managing Avi Load Balancer in VMware Cloud Foundation

VMware® Avi™ Load Balancer (formerly known as NSX Advanced Load Balancer) allows you to implement centrally-managed distributed load balancing for your application workloads within VMware Cloud Foundation and configure enterprise grade load-balancing, global server load balancing, application security, and container ingress services.

Starting with VMware Cloud Foundation 5.2, you can use SDDC Manager to deploy Avi Load Balancer as a high availability cluster of three VMware® Avi™ Controller instances, each running on a separate VM.

NOTE

Previous version of VMware Cloud Foundation support Avi Load Balancer, but do not deploy or manage the Avi Controller Cluster.

The Avi Controller cluster functions as the control plane and stores and manages all policies related to services and management. All Avi Controllers are deployed in the management domain, even when the Avi Load Balancer is deployed in a VI workload domain.

When you deploy Avi Load Balancer in a workload domain, it is associated with the workload domain's NSX Manager.

NOTE

VMware Cloud Foundation 5.2 does not support deploying Avi Load Balancer on a workload domain that shares its NSX Manager with another workload domain.

VMware Cloud Foundation does not deploy or manage the Service Engine VMs (SEs) that function as the data plane. After deploying the Avi Controller cluster, you can use the Avi Load Balancer UI/API, VMware Aria Automation, or Avi Kubernetes Operator to deploy virtual services for an application, which creates the required Service Engine virtual machines. Service Engines (SEs) are deployed in the workload domain in which the Avi Load Balancer is providing load

balancing services. All SEs deployed in a VI workload domain are managed by the Avi Controller that is part of the Avi Load Balancer deployment that is associated with the corresponding NSX instance managing the VI workload domain.

Other important considerations:

- VMware Cloud Foundation does not manage license updates for Avi Load Balancer.
- VMware Cloud Foundation does not manage backing up of Avi Load Balancer configuration database. See the [VMware Avi Load Balancer Documentation](#) for information about configuring scheduled and on-demand backups.
- VMware Cloud Foundation does not manage upgrading Avi Controller Cluster. See the [VMware Avi Load Balancer Documentation](#) for information about upgrading.
- The lifecycle of the Avi Service Engines is managed by each Avi Controller Cluster. You perform updates and upgrades in the Avi Load Balancer web interface, which has checks in place to ensure that you can only upgrade to supported versions.
- If you upgraded from an earlier version of VMware Cloud Foundation and had deployed Avi Load Balancer, SDDC Manager will not be aware of or manage that Avi Load Balancer. You can use SDDC Manager to deploy additional Avi Load Balancers in such an environment.
- In order to use Avi Load Balancer for load balancing services in a vSphere IaaS Control Plane environment, the Avi Load Balancer must be registered with the NSX Manager. See [Registering an Avi Load Balancer cluster with an NSX Manager instance](#).

For more information about how to use and manage Avi Load Balancer see:

- [VMware Avi Load Balancer Documentation](#)
- The [Advanced Load Balancing for VMware Cloud Foundation](#) validated solution

Limitations of Avi Load Balancer in VMware Cloud Foundation

Avi Load Balancer is not supported in workload domains with a shared NSX Manager instance.

Deploy Avi Load Balancer for a Workload Domain

You can deploy Avi Load Balancer for workload domains that do not share their NSX Manager with any other workload domains.

Download the install bundle for a supported version of NSX Advanced Load Balancer. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).

Avi Load Balancer was formerly known as NSX Advanced Load Balancer. The SDDC Manager UI still refers to NSX Advanced Load Balancer.

You cannot deploy Avi Load Balancer on a workload domain that shares its NSX Manager with another workload domain.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Select **Actions** > **Deploy NSX Advanced Load Balancer**.
4. Select the NSX Advanced Load Balancer version and click **Next**.

Select the version of NSX Advanced Load Balancer to deploy

Select version

NEXT

Select ▾

Select

22.1.6

5. Select the appliance size and click **Next**.

Select the size of the NSX Advanced Load Balancer instance to be deployed

Select Size Medium ▼
 Select a deployment size to see resource usage against availability

Resource Availability

Resource	Available in Management Domain	Resource reservation needed
Memory	246.62 GB	96 GB (38.93 % of available)
CPU	103.1 GHz	48 GHz (46.56 % of available)
Disk Space	3.38 TB	768 GB (22.19 % of available)

NEXT

Make sure that the management domain has enough resources for the selected size.

6. Enter the settings for the NSX Advanced Load Balancer Controller cluster and click **Next**.

Administrator password	Enter an administrator password. Although individual VMware Cloud Foundation accounts support different password requirements, it is recommended that you set passwords following a common set of requirements across all accounts: <ul style="list-style-type: none"> • Minimum length: 12 • Maximum length: 20 • At least one lowercase letter, one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ * • Must NOT include: <ul style="list-style-type: none"> – A dictionary word – A palindrome – More than four monotonic character sequences – Three of the same consecutive characters
Node 1 IP Address	Enter the IP address of the first Avi Controller.
Node 2 IP Address	Enter the IP address of the second Avi Controller.
Node 3 IP Address	Enter the IP address of the third Avi Controller.
Cluster VIP	Enter the Avi Load Balancer Controller cluster IP address. The Avi Load Balancer Controller cluster IP address is a single IP address shared by the Avi Controllers within the cluster. It is the address to which the web interface, CLI commands, and REST API calls are directed. As a best practice, to access the Avi Controller, you must log in to the cluster IP address instead of the IP addresses of individual Avi Controller nodes.
Cluster FQDN	Enter the Avi Controller cluster FQDN.

Table continued on next page

Continued from previous page

	<p>NOTE When creating a service account for the NSX Advanced Load Balancer Controller cluster, VMware Cloud Foundation 5.2 combines the Avi Load Balancer VIP host name and the NSX Manager VIP host name to create the account, <code>svc-<alb hostname>-<nsx hostname></code>. The total characters cannot exceed 32. VCF 5.2.1 automatically truncates the service account name to avoid deployment failures based on account name length.</p>
Cluster Name	Enter a name for the Avi Controller cluster.

7. Click **Start Deployment**.

You can monitor the deployment in the Tasks panel.

Task	Subtask	Task Status	Last Occurrence
Adding NSX Advanced Load Balancer Cluster nsx-alb	Deploy NSX Advanced Load Balancer on NSX. This includes OVA upload, VM creation and wait for cluster HA.	<div style="width: 42%;"><div style="background-color: #0070c0; height: 10px;"></div></div> 42%	5/11/24, 1:05 AM
Download BUNDLE - NSX_ALB:22.16-23390967		Successful	5/10/24, 11:58 PM

After the Avi Load Balancer Controller cluster deploys successfully, you can access the web interface from the **Services** tab for the workload domain by clicking the NSX Advanced Load Balancer link.

MANAGEMENT ✔ ACTIVE Version : 5.2.0.0

Summary Services Updates Update History

VMware Cloud Foundation Components

Component
vCenter Server vcenter-1.vrack.vsphere.local ↗
NSX Cluster vip-nsx-mgmt.vrack.vsphere.local ↗
NSX Advanced Load Balancer vipa.vrack.vsphere.local ↗

You can manage the Avi Load Balancer Controller cluster administrator password and certificate using the SDDC Manager UI.

- [Managing Passwords in VMware Cloud Foundation](#)
- [Managing Certificates in VMware Cloud Foundation](#)

Remove Avi Load Balancer from a Workload Domain

If you deployed an Avi Load Balancer to a workload domain and you no longer need it, you can remove it from the workload domain. If a workload domain includes an Avi Load Balancer, you cannot delete the workload domain until you remove the Avi Load Balancer.

The Avi Load Balancer must not be in use.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. In the **Workload Domains** page, click a domain name in the Domain column.
3. Select **Actions** › **Remove NSX Advanced Load Balancer**.

Add Cluster
 Add Edge Cluster
 Add AVNs
 Update Licenses
 Rename Domain
 Remove NSX Advanced Load Balancer

4. Click **Remove** to confirm.

Remove NSX Advanced Load Balancer



Are you sure you want to remove NSX Advanced Load Balancer from this workload domain?

CANCEL

REMOVE

Deploying Application Virtual Networks in VMware Cloud Foundation

Before you can deploy VMware Aria Suite components or implement the Identity and Access Management for VMware Cloud Foundation validated solution, you must deploy Application Virtual Networks in the management domain.

An Application Virtual Network (AVN) is a software-defined networking concept based on NSX that allows the hosting of management applications on NSX segments. In NSX, segments are virtual layer-2 domains.

You can create overlay-backed NSX segments or VLAN-backed NSX segments. Both options create two NSX segments (Region-A and X-Region) on the NSX Edge cluster deployed in the default management vSphere cluster. Those NSX segments are used when you deploy the VMware Aria Suite products. Region-A segments are local instance NSX segments and X-Region segments are cross-instance NSX segments.

IMPORTANT

You cannot create AVNs if the NSX for the management domain is part of an NSX Federation.

Overlay-Backed NSX Segments

Overlay-backed segments provide flexibility for workload placement by removing the dependence on traditional data center networks. Using overlay-backed segments improves the security and mobility of management applications and reduces the integration effort with existing networks. Overlay-backed segments are created in an overlay transport zone.

In an overlay-backed segment, traffic between two VMs on different hosts but attached to the same overlay segment have their layer-2 traffic carried by a tunnel between the hosts. NSX instantiates and maintains this IP tunnel without the need for any segment-specific configuration in the physical infrastructure. As a result, the virtual network infrastructure is decoupled from the physical network infrastructure. That is, you can create segments dynamically without any configuration of the physical network infrastructure.

VLAN-Backed NSX Segments

VLAN-backed segments leverage the physical data center networks to isolate management applications, while still taking advantage of NSX to manage these networks. VLAN-backed network segments ensure the security of management applications without requiring support for overlay networking. VLAN-backed segments are created in a VLAN transport zone.

A VLAN-backed segment is a layer-2 broadcast domain that is implemented as a traditional VLAN in the physical infrastructure. This means that traffic between two VMs on two different hosts but attached to the same VLAN-backed segment is carried over a VLAN between the two hosts. The resulting constraint is that you must provision an appropriate VLAN in the physical infrastructure for those two VMs to communicate at layer-2 over a VLAN-backed segment.

VMware Aria Suite Components and NSX Segments

When you deploy the VMware Aria Suite components, they use the NSX segments that you created.

VMware Aria Suite Component	NSX Segment
VMware Aria Operations for Logs	Region-A
VMware Aria Operations Manager	X-Region
Workspace ONE Access	X-Region
VMware Aria Automation	X-Region
VMware Aria Suite Lifecycle	X-Region

Identity and Access Management for VMware Cloud Foundation

See [Identity and Access Management for VMware Cloud Foundation](#) for more information about how that validated solution uses Application Virtual Networks.

Deploy Overlay-Backed NSX Segments

Create overlay-backed NSX segments, also known as Application Virtual Networks (AVNs), for use with VMware Aria Suite components.

Create an NSX Edge cluster for Application Virtual Networks, using the recommended settings, in the default management vSphere cluster. See [Deploy an NSX Edge Cluster](#).

This procedure describes creating overlay-backed NSX segments. If you want to create VLAN-backed NSX segments instead, see [Deploy VLAN-Backed NSX Segments](#).

1. In the navigation page, click **Inventory** › **Workload Domains**.
2. Click on the management domain.

3. Select **Actions** > **Add AVNs**.
4. Select **Overlay-backed network segment** and click **Next**.
5. Select an NSX Edge cluster and a Tier-1 gateway.
6. Enter information for each of the NSX segments (Region-A and X-Region):

Option	Description
Name	Enter a name for the NSX segment. For example, <code>Mgmt-RegionA01</code> .
Subnet	Enter a subnet for the NSX segment.
Subnet mask	Enter a subnet mask for the NSX segment.
Gateway	Enter a gateway for the NSX segment.
MTU	Enter an MTU for the NSX segment.

7. Click **Validate Settings** and then click **Next**.
If validation does not succeed, verify and update the information you entered for the NSX segments and click **Validate Settings** again.
8. Review the settings and click **Finish**.

Example Network Topology for Overlay-Backed NSX Segments



Deploy VLAN-Backed NSX Segments

Create VLAN-backed NSX segments, also known as Application Virtual Networks (AVNs), for use with VMware Aria Suite components.

Create an NSX Edge cluster for Application Virtual Networks, using the recommended settings, in the default management vSphere cluster. See [Deploy an NSX Edge Cluster](#).

You must have an available VLAN ID for each NSX segment.

This procedure describes creating VLAN-backed NSX segments. If you want to create overlay-backed NSX segments instead, see [Deploy Overlay-Backed NSX Segments](#).

1. In the navigation page, click **Inventory** › **Workload Domains**.
2. Click on the management domain.
3. Select **Actions** › **Add AVNs**.
4. Select **VLAN-backed network segment** and click **Next**.
5. Select an NSX Edge cluster.
6. Enter information for each of the NSX segments (Region-A and X-Region):

Option	Description
Name	Enter a name for the NSX segment. For example, Mgmt-RegionA01.
Subnet	Enter a subnet for the NSX segment.
Gateway	Enter a gateway for the NSX segment.
MTU	Enter an MTU for the NSX segment.
VLAN ID	Enter the VLAN ID for the NSX segment.

7. Click **Validate Settings** and then click **Next**.

If validation does not succeed, verify and update the information you entered for the NSX segments and click **Validate Settings** again.

8. Review the settings and click **Finish**.

Example Network Topology for VLAN-Backed NSX Segments



VMware Cloud Foundation with VMware Tanzu

VMware Cloud Foundation™ with VMware Tanzu™ enables you to deploy and operate the compute, networking, and storage infrastructure for vSphere IaaS Control Plane workloads. vSphere IaaS Control Plane transforms vSphere to a platform for running Kubernetes workloads natively on the hypervisor layer.

When enabled on a vSphere cluster, vSphere IaaS Control Plane provides the capability to run Kubernetes workloads directly on ESXi hosts and to create upstream Kubernetes clusters within dedicated resource pools. vSphere IaaS Control Plane can also be enabled on the management domain default cluster.

NOTE

Starting with vSphere 8.0 Update 3, vSphere with Tanzu was renamed to vSphere IaaS Control Plane.

You validate the underlying infrastructure for vSphere IaaS Control Plane from the SDDC Manager UI and then complete the deployment in the vSphere Client. The SDDC Manager UI refers to the vSphere IaaS Control Plane functionality as Kubernetes - Workload Management.

The [Developer Ready Infrastructure for VMware Cloud Foundation](#) validated solution provides design, implementation, and operational guidance for a workload domain that runs vSphere with Tanzu workloads in the Software-Defined Data Center (SDDC).

For more information about vSphere IaaS Control Plane, see [What Is vSphere IaaS Control Plane?](#)

Enable Workload Management

With Workload Management, you validate the underlying infrastructure for vSphere IaaS Control Plane. You then complete the deployment using the vSphere Client.

- A VI workload domain must be deployed.

NOTE

If you deployed VMware Cloud Foundation with a consolidated architecture, you can enable Workload Management on the management domain.

- An Workload Management ready NSX Edge cluster must be deployed on the workload domain. You must select Workload Management on the Use Case page of the Add Edge Cluster wizard. See step 6 in [Deploy an NSX Edge Cluster](#).
- All hosts in the vSphere cluster for which you enable Workload Management must be licensed for vSphere IaaS Control Plane.
- Workload Management requires a vSphere cluster with a minimum of three ESXi hosts.
- The following IP address subnets must be defined:
 - A non-routable subnet for pod networking, minimum of a /22 subnet.
 - A non-routable subnet for Service IP addresses, minimum of a /24 subnet
 - A routable subnet for ingress, minimum of a /27 subnet
 - A routable subnet for egress, minimum of a /27 subnet
- In order to use Avi Load Balancer for load balancing services in a vSphere IaaS Control Plane environment, the Avi Load Balancer must be registered with the NSX Manager. See [Registering an Avi Load Balancer cluster with an NSX Manager instance](#).

1. In the navigation pane, click **Solutions**.
2. In the Kubernetes - Workload Management section, click **Deploy**.

Solutions

Kubernetes - Workload Management

With Workload Management, you can deploy and configure the compute, networking, and storage infrastructure for vSphere with Kubernetes.

No Workload Management solution has been created.

[LEARN MORE](#) [VIEW DETAILS](#) [DEPLOY](#)

3. Review the Workload Management prerequisites, click **Select All**, and click **Begin**.
4. Select the workload domain associated with the vSphere cluster where you want to enable Workload Management. The Workload Domain drop-down menu displays all Workload Management ready workload domains, including the management domain.

vSphere clusters in the selected workload domain that are compatible with Workload Management are displayed in the Compatible section. Incompatible clusters are displayed in the Incompatible section, along with the reason for the incompatibility. If you want to get an incompatible cluster to a usable state, you can exit the Workload Management deployment wizard while you resolve the issue.

5. From the list of compatible clusters on the workload domain, select the cluster where you want to enable Workload Management and click **Next**.
6. On the Validation page, wait for validation to complete successfully and click **Next**.

The following validations are performed.

- vCenter Server validation (vCenter Server credentials, vSphere cluster object, and version)
- Network validation (NSX Manager credentials and version)
- Compatibility validation (vSphere cluster and content library)

7. On the Review page, review your selections and click **Complete in vSphere**. You are automatically redirected to the vSphere Client.

Follow the deployment wizard within the vSphere Client to complete the Workload Management deployment and configuration steps.

View Workload Management Cluster Details

The Workload Management page displays clusters with Workload Management. The status of each cluster, number of hosts in the cluster, and associated workload domain is also displayed.

1. In the navigation pane, click **Solutions**.
2. In the Kubernetes - Workload Management section, click **View Details**.
3. Click vSphere Workload Management Clusters to see cluster details in vSphere.

Update Workload Management License

Once you enable Workload Management on a cluster, you must assign a Tanzu edition license to the cluster before the evaluation license expires.

You must have added a VMware Tanzu license key to the Cloud Foundation license inventory. See [Add a Component License Key in the SDDC Manager UI](#).

1. In the navigation pane, click **Solutions**.
2. Click the dots to the left of the cluster for which you want to update the license and click **Update Workload Management license**.
3. Select the appropriate license and click **Apply**.
After the license update processing is completed, the Workload Management page is displayed. The task panel displays the licensing task and its status.

VMware Aria Suite Lifecycle in VMware Cloud Foundation mode

When you deploy VMware Aria Suite Lifecycle from the SDDC Manager UI, VMware Cloud Foundation mode is enabled in VMware Aria Suite Lifecycle, and the behavior of VMware Aria Suite Lifecycle is aligned with the VMware Cloud Foundation architecture.

VMware Aria Suite Lifecycle in VMware Cloud Foundation mode introduces the following features:

- Automatic load balancer configuration. Load balancer preparation and configuration are no longer a prerequisite when you use VMware Aria Suite Lifecycle to deploy or perform a cluster expansion on Workspace ONE Access, VMware Aria Operations, or VMware Aria Automation. Load balancer preparation and configuration take place as part of the deploy or expand operation.
- Automatic infrastructure selection in the VMware Aria Suite Lifecycle deployment wizards. When you deploy a VMware Aria Suite product through VMware Aria Suite Lifecycle, infrastructure objects such as clusters and networks are pre-populated. They are fixed and cannot be changed to ensure alignment with the VMware Cloud Foundation architecture.
- Cluster deployment for a new environment. You can deploy VMware Aria Operations for Logs, VMware Aria Operations, or VMware Aria Automation in clusters. You can deploy Workspace ONE Access either as a cluster or a single node. If you deploy Workspace ONE Access as a single node, you can expand it to a cluster later.
- Consistent Bill Of Materials (BOM). VMware Aria Suite Lifecycle in VMware Cloud Foundation mode only displays product versions that are compatible with VMware Cloud Foundation to ensure product interoperability.
- Inventory synchronization between VMware Aria Suite Lifecycle and SDDC Manager. VMware Aria Suite Lifecycle can detect changes made to VMware Aria Suite products and update its inventory through inventory synchronization. When VMware Cloud Foundation mode is enabled in VMware Aria Suite Lifecycle, inventory synchronization in VMware Aria Suite Lifecycle also updates SDDC Manager's inventory to get in sync with the current state of the system.
- Product versions. VMware Cloud Foundation supports flexible VMware Aria Suite upgrades. You can upgrade VMware Aria Suite products as new versions become available in VMware Aria Suite Lifecycle. VMware Aria Suite Lifecycle will only allow upgrades to compatible and supported versions of VMware Aria Suite products.
- Resource pool and advanced properties. The resources in the Resource Pools under the Infrastructure Details are blocked by the VMware Aria Suite Lifecycle UI, so that the VMware Cloud Foundation topology does not change. Similarly, the Advanced Properties are also blocked for all products except for Remote Collectors. VMware Aria Suite Lifecycle also auto-populates infrastructure and network properties by calling VMware Cloud Foundation deployment API.
- Federal Information Processing Standard (FIPS) compliance.
- Watermark.

VMware Aria Suite Lifecycle Implementation

You deploy VMware Aria Suite Lifecycle in VMware Cloud Foundation mode by using SDDC Manager. After that, you perform the necessary post-deployment configurations.

- Download the VMware Software Install Bundle for VMware Aria Suite Lifecycle from the VMware Depot to the local bundle repository. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Allocate an IP address for the VMware Aria Suite Lifecycle virtual appliance on the cross-instance NSX segment and prepare both forward (A) and reverse (PTR) DNS records.
- Allocate an IP address for the NSX standalone Tier-1 Gateway on the cross-instance NSX segment. This address is used for the service interface of the standalone NSX Tier 1 Gateway created during the deployment. The Tier 1 Gateway is used for load-balancing of specific VMware Aria Suite products and Workspace ONE Access.
- Ensure you have enough storage capacity:
 - Required storage: 178 GB
 - Virtual disk provisioning: Thin
- Verify that the management domain vCenter Server is operational.
- Verify that NSX Manager is operational.
- Verify the **Prerequisite Checklist** sheet in the *Planning and Preparation Workbook*.

By default, VMware Cloud Foundation uses NSX to create NSX segments and deploys VMware Aria Suite Lifecycle and the VMware Aria Suite products to these NSX segments. Starting with VMware Cloud Foundation 4.3, NSX segments are no longer configured during the management domain bring-up process, but instead are configured using the SDDC Manager UI. The new process offers the choice of using either overlay-backed or VLAN-backed segments. See [Deploying Application Virtual Networks in VMware Cloud Foundation](#).

VMware Aria Suite Lifecycle runs in VMware Cloud Foundation mode, the integration ensures awareness between the two components. You launch the deployment of VMware Aria Suite products from the SDDC Manager UI and are redirected to the VMware Aria Suite Lifecycle UI where you complete the deployment process.

Deploy VMware Aria Suite Lifecycle

You deploy the VMware Aria Suite Lifecycle in VMware Cloud Foundation mode by using the SDDC Manager UI.

1. In the navigation pane, click **Administration** > **VMware Aria Suite**.
2. Click **Deploy**.
3. Review and verify the prerequisites.
 - Click each prerequisite check box and then click **Begin**.
4. On the **Network Settings** page, review the settings and click **Next**.
5. On the **Virtual Appliance Settings** page, enter the settings and click **Next**.

Setting	Description
Virtual Appliance: FQDN	<p>The FQDN for the VMware Aria Suite Lifecycle virtual appliance.</p> <p>NOTE The reverse (PTR) DNS record of this fully qualified domain name is used as the IP address for the virtual appliance.</p>

Table continued on next page

Continued from previous page

Setting	Description
NSX Tier 1 Gateway: IP Address	<p>A free IP Address within the cross-instance virtual network segment.</p> <p>NOTE Used to create a service interface on the NSX Tier 1 Gateway, where VMware Cloud Foundation automatically configures the load-balancer for the VMware Aria Suite.</p>
System Administrator	<p>Create and confirm the password for the VMware Aria Suite Lifecycle administrator account, vcfadmin@local. The password created is the credential that allows SDDC Manager to connect to VMware Aria Suite Lifecycle.</p> <p>NOTE When VMware Aria Suite Lifecycle is deployed by SDDC Manager it is enabled for VMware Cloud Foundation mode. As a result, the administrator account for is vcfadmin@local instead of admin@local.</p>
SSH Root Account	<p>Create and confirm a password for the VMware Aria Suite Lifecycle virtual appliance root account.</p>

6. On the **Review Summary** page, review the installation configuration settings and click **Finish**. SDDC Manager validates the values and starts the deployment.

The VMware Aria Suite page displays the following message: `Deployment in progress`.

If the deployment fails, this page displays a deployment status of `Deployment failed`. In this case, you can click **Restart Task** or **Rollback**.

7. **(Optional)** To view details about the individual deployment tasks, in the **Tasks** panel at the bottom, click each task.

Replace the Certificate of the VMware Aria Suite Lifecycle Instance

To establish a trusted connection to VMware Aria Suite Lifecycle, you replace the SSL certificate on the appliance by using the SDDC Manager UI.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. On the **Workload Domain** page, from the table, in the domain column click the management domain.
3. On the domain summary page, click the **Certificates** tab.
4. From the table, select the check box for the VMware Aria Suite Lifecycle resource type, and click **Generate CSRs**.
5. On the **Details** page, enter the following settings and click **Next**.

Settings	Description
Algorithm	Select the key algorithm for the certificate.
Key Size	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
Email	Optionally, enter a contact email address.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization Name	Type the name under which your company is known. The listed organization must be the legal registrant of the domain name in the certificate request.
Locality	Type the city or locality where your company is legally registered.
State	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- On the **Subject Alternative Name** page, leave the default SAN and click **Next**.
- On the **Summary** page, click **Generate CSRs**.
- After the successful return of the operation, click **Generate signed certificates**.
- In the **Generate Certificates** dialog box, from the **Select Certificate Authority** drop-down menu, select **Microsoft**.
- Click **Generate certificates**.
- After the successful return of the operation, click **Install certificates**.
Wait for the successful return of the operation.

Configure Data Center and vCenter Server in VMware Aria Suite Lifecycle

Before you can create a global environment for product deployments, you must add a cross-instance data center and the associated management domain vCenter Server to VMware Aria Suite Lifecycle.

You add the cross-instance data center, and the associated management domain vCenter Server for the deployment of the global components, such as the clustered Workspace ONE Access.

- In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
- On the **My Services** page, click **Lifecycle Operations**.
- In the navigation pane, click **Datacenters**.
- Click **Add datacenter**, enter the values for the global data center, and click **Save**.

Setting	Value
Datacenter name	Name for cross-instance datacenter
Use custom location	Deactivated
Location	Location of datacenter

- Add the management domain vCenter Server to the global data center.

- a) On the **Datcenters** page, expand the global data center and click **Add vCenter**.
- b) Enter the management domain vCenter Server information and click **Validate**.

Setting	Value
vCenter name	Enter a name for the vCenter Server
vCenter FQDN	Enter the FQDN of the vCenter Server
vCenter credentials	Select the <code><management_vcenter_name>-<uid></code> credential. For example: vcenter-1-35214fac-caeb-4062-a184-350344e30c7f.
vCenter type	Management

6. After the successful vCenter Server validation, click **Save**.
7. In the navigation pane, click **Requests** and verify that the state of the **vCenter data collection request** is Completed.

Workspace ONE Access Implementation

Workspace ONE Access provides identity and access management services for the VMware Aria Suite of products. You use VMware Aria Suite Lifecycle to deploy a Workspace ONE Access instance. You then perform the necessary post-deployment configurations and customization. VMware Cloud Foundation supports both standard and clustered Workspace ONE Access deployments.

- Download the installation binary directly from VMware Aria Suite Lifecycle. See "Configure Product Binaries" in the *VMware Aria Suite Lifecycle Installation, Upgrade, and Management Guide* for the version of [VMware Aria Suite Lifecycle](#) listed in the VMware Cloud Foundation BOM.
- Allocate IP addresses:

Standard Deployment	Clustered Deployment
One IP address from the cross-instance NSX segment and prepare both forward (A) and reverse (PTR) DNS records.	Five IP addresses from the cross-instance NSX segment and prepare both forward (A) and reverse (PTR) DNS records. <ul style="list-style-type: none"> • Three IP addresses for the clustered Workspace ONE Access instance. • One IP address for the embedded Postgres database for the Workspace ONE Access instance. • One IP address for the NSX external load balancer virtual server for clustered Workspace ONE Access instance.

- Ensure you have enough storage capacity:
 - Required storage per node: 100 GB
 - Virtual disk provisioning: Thin
- Verify that the management domain vCenter Server is operational.
- Verify that the cross-instance NSX segment is available.
- Verify that the NSX Manager is operational.
- Verify the **Prerequisite Checklist** sheet in the *Planning and Preparation Workbook*.
- Verify that required Active Directory bind service account is created.

Verify that required Active Directory security groups are created.

- Download the [CertGenVVS](#) tool and generate the signed certificate for the Workspace ONE Access instance. See [KB 85527](#).

Import the Workspace ONE Access Certificate to VMware Aria Suite Lifecycle

To prepare VMware Aria Suite Lifecycle for deploying Workspace ONE Access, you must generate an SSL certificate using the PowerShell module for VMware Validated Solutions and add the certificate to the VMware Aria Suite Lifecycle locker.

- Verify that a Microsoft Certificate Authority is available for the environment.
- Install the [PowerShell module for VMware Validated Solutions](#) together with the supporting modules to request an SSL certificate from your Microsoft Certificate Authority.
- Verify that you have OpenSSL 3.0 or later installed on the system that will run the PowerShell module. The [OpenSSL Wiki](#) has a list of third-party pre-compiled binaries for Microsoft Windows.

This procedure uses the PowerShell Module for VMware Validated Solutions to generate the required certificates from a Microsoft Active Directory Certificate Services. However, the module also supports generating certificate signing requests (CSRs) for third party certificate authorities for import to the VMware Aria Suite Lifecycle locker.

1. Generate an SSL certificate using the PowerShell module for VMware Validated Solutions.
 - a) Start PowerShell.
 - b) Replace the sample values in the variables below and run the commands in the PowerShell console.

```
$commonName = "xint-idm01.rainpole.io"

$subjectAltNames = "xint-idm01.rainpole.io, xint-idm01a.rainpole.io, xint-
idm01b.rainpole.io, xint-cidm01c.rainpole.io"

$encryptionKeySize = 2048

$certificateExpiryDays = 730

$orgName = "rainpole"

$orgUnitName = "Platform Engineering"

$orgLocalityName = "San Francisco"

$orgStateName = "California"

$orgCountryCode = "US"

$caType = "msca"

$caFqdn = "rpl-ad01.rainpole.io"

$caUsername = "Administrator"

$caPassword = "VMw@re1!"

$caTemplate = "VMware"

$outputPath = ".\certificates\"

$csrFilePath = Join-Path $outputPath "$commonName.csr"

$keyFilePath = Join-Path $outputPath "$commonName.key"

$certFilePath = Join-Path $outputPath "$commonName.crt"
```



```
$rootCaFilePath = Join-Path $outputPath "$caFqdn-rootCa.pem"
```

c) Perform the configuration by running the command in the PowerShell console.

```
Invoke-GeneratePrivateKeyAndCsr -outDirPath $outputPath -commonName $commonName
-subjectAlternativeNames $subjectAltNames -keySize $encryptionKeySize
-expireDays $certificateExpiryDays -organization $orgName -organizationUnit
$orgUnitName -locality $orgLocalityName -state $orgStateName -country
$orgCountryCode
```

```
Invoke-RequestSignedCertificate -caFqdn $caFqdn -csrFilePath $csrFilePath
-outDirPath $outputPath -certificateAuthority $caType -username $caUsername
-password $caPassword -certificateTemplate $caTemplate -getCArootCert
```

```
Invoke-GenerateChainPem -outDirPath $outputPath -keyFilePath $keyFilePath
-crtFilePath $crtFilePath -rootCaFilePath $rootCaFilePath
```

2. Add the generated SSL certificate to the VMware Aria Suite Lifecycle locker.

- a) Log in to VMware Aria Suite Lifecycle at https://<aria_suite_lifecycle_fqdn> as `vcfadmin@local`.
- b) On the **My services** page, click **Locker**.
- c) In the navigation pane, click **Certificates**.
- d) On the **Certificates** page, click **Import**.
- e) On the **Import certificate** page, enter a name for the Workspace ONE Access certificate according to your *VMware Cloud Foundation Planning and Preparation Workbook*.
- f) Click **Browse file**, navigate to the Workspace ONE Access certificate file (.pem), and click **Open**.
- g) On the **Import certificate** page, click **Import**.

Add Workspace ONE Access Passwords to VMware Aria Suite Lifecycle

To enable life cycle management and configuration management, you set the passwords for the VMware Aria Suite Lifecycle cross-instance environment administrator account and for the Workspace ONE Access administrator and configuration administrator accounts.

You add the following passwords for the corresponding local administrative accounts.

Setting	Value for Global Environment Administrator	Value for Local Administrator	Value for Local Configuration Administrator	Value for Appliance Root User
Password alias	global-env-admin	xint-wsa-admin	xint-wsa-configadmin	xint-wsa-root
Password	<code>global_env_admin_password</code>	<code>xint_wsa_admin_password</code>	<code>xint_wsa_configadmin_password</code>	<code>xint_wsa_root_password</code>
Confirm password	<code>global_env_admin_password</code>	<code>xint-wsa_admin_password</code>	<code>xint_wsa_configadmin_password</code>	<code>xint_wsa_root_password</code>
Password description	VMware Aria Suite Lifecycle global environment default password Used for Workspace ONE Access appliance sshuser.	Workspace ONE Access administrator	Workspace ONE Access configuration administrator	Workspace ONE Access root user

NOTE

You do not need to provide a user name when adding passwords. You can leave the **User Name** field blank when configuring settings.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Locker**.
3. In the navigation pane, click **Passwords**.
4. On the **Passwords** page, click **Add**.
5. On the **Add password** page, configure the settings and click **Add**.
6. Repeat this procedure for all the remaining credentials.

Deploy a Standard Workspace ONE Access Instance Using VMware Aria Suite Lifecycle

To provide identity and access management services to the cross-instance SDDC components, you create a global environment in VMware Aria Suite Lifecycle in which you deploy a standard Workspace ONE Access instance.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Lifecycle Operations**.
3. On the **Dashboard** page, click **Create environment**.
4. On the **Create environment** page, configure the settings and click **Next**.

Setting	Value
Install Identity Manager	Selected
Default password	global-env-admin
Datacenter	Select the cross-instance datacenter.
JSON configuration	Deactivated
Join the VMware customer experience improvement program	Selected

5. On the **Select product** page, select the check box for **VMware Identity Manager**, configure these values, and click **Next**.

Setting	Value
Installation type	New install
Version	Select a version. VMware Aria Suite Lifecycle will only display supported versions.
Deployment type	Standard

6. On the **Accept license agreements** page, scroll to the bottom and accept the license agreement, and then click **Next**.
7. On the **Certificate** page, from the **Select certificate** drop-down menu, select the *Workspace One Access* certificate, and click **Next**.
8. On the **Infrastructure** page, verify and accept the default settings, and click **Next**.
9. On the **Network** page, verify and accept the default settings, and click **Next**.
10. On the **Products** page, configure the deployment properties of Workspace ONE Access and click **Next**.

- a) In the **Product properties** section, configure the settings.

Setting	Value
Certificate	<i>Workspace One Access</i>
Node size	Medium (VMware Aria Automation recommended size)
Admin password	Select the <i>xint-wsa-admin</i>
Default configuration admin email	Enter a default email.
Default configuration admin user name	<i>configadmin</i>
Default configuration admin password	Select the <i>xint-wsa-configadmin</i>
Sync group members	Selected

- b) In the **Components** section, configure the primary node.

Setting	Value for vidm-primary
VM Name	Enter a VM Name for vidm-primary.
FQDN	Enter the FQDN for vidm-primary
IP address	Enter the IP Address for vidm-primary.

- c) Click advanced configuration and click **Select Root Password**.

- d) Select *xint-wsa-root* and click **Save**.

11. On the **Precheck** page, click **Run precheck**.
12. On the **Manual validations** page, select the **I took care of the manual steps above and am ready to proceed** check box and click **Run precheck**.
13. Review the validation report, remediate any errors, and click **Re-run precheck**.
14. Wait for all prechecks to complete with *Passed* messages and click **Next**.
15. On the **Summary** page, review the configuration details. To back up the deployment configuration, click **Export configuration**.
16. To start the deployment, click **Submit**.
The **Request details** page displays the progress of deployment.
17. Monitor the steps of the deployment graph until all stages become *Completed*.

Deploy Clustered Workspace ONE Access Instance Using VMware Aria Suite Lifecycle

To provide identity and access management services to the cross-instance SDDC components, you create a global environment in VMware Aria Suite Lifecycle in which you deploy a 3-node clustered Workspace ONE Access instance.

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Lifecycle Operations**.
3. On the **Dashboard** page, click **Create environment**.
4. On the **Create environment** page, configure the settings and click **Next**.

Setting	Value
Install Identity Manager	Selected
Default password	global-env-admin
Datacenter	Select the cross-instance datacenter.
JSON configuration	Deactivated
Join the VMware customer experience improvement program	Selected

5. On the **Select product** page, select the check box for **VMware Identity Manager**, configure these values, and click **Next**.

Setting	Value
Installation type	New install
Version	Select a version. VMware Aria Suite Lifecycle will only display supported versions.
Deployment type	Cluster

6. On the **Accept license agreements** page, scroll to the bottom and accept the license agreement, and then click **Next**.
7. On the **Certificate** page, from the **Select certificate** drop-down menu, select the *Clustered Workspace One Certificate*, and click **Next**.
8. On the **Infrastructure** page, verify and accept the default settings, and click **Next**.
9. On the **Network** page, verify and accept the default settings, and click **Next**.
10. On the **Products** page, configure the deployment properties of clustered Workspace ONE Access and click **Next**.
- a) In the **Product properties** section, configure the settings.

Setting	Value
Certificate	<i>Workspace One Access</i>
Node size	Medium (VMware Aria Automation recommended size)
Admin password	Select the <i>xint-wsa-admin</i>
Default configuration admin email	Enter a default email.
Default configuration admin user name	<i>configadmin</i>
Default configuration admin password	Select the <i>xint-wsa-configadmin</i>
Sync group members	Selected

- b) In the **Cluster Virtual IP** section, click **Add Load Balancer** and configure its settings.

Setting	Value
Controller Type	VMware Cloud Foundation managed NSX-T

Table continued on next page

Continued from previous page

Setting	Value
Load Balancer IP	Use the IP address from your <i>VMware Cloud Foundation Planning and Preparation Workbook</i> .
Load Balancer FQDN	Use the FQDN from your <i>VMware Cloud Foundation Planning and Preparation Workbook</i> .

c) In the **Cluster VIP FQDN** section, configure the settings.

Setting	Value
Controller Type	Select VMware Cloud Foundation managed NSX-T from the drop-down menu.
FQDN	Select the Load Balancer FQDN from the drop-down menu.
Locker certificate	Clustered Workspace ONE Access Certificate
Database IP address	Enter the IP address for the embedded Postgres database. NOTE The IP address must be a valid IP address for the cross-instance NSX segment.

d) In the **Components** section, configure the three cluster node.

Setting	Value for vidm-primary	Value for vidm-secondary-1	Value for vidm-secondary-2
VM Name	Enter a VM Name for vidm-primary.	Enter a VM Name for vidm-secondary-1.	Enter a VM Name for vidm-secondary-2.
FQDN	Enter the FQDN for vidm-primary	Enter the FQDN for vidm-secondary-1.	Enter the FQDN for vidm-secondary-2.
IP address	Enter the IP Address for vidm-primary.	Enter the IP Address for vidm-secondary-1.	Enter the IP Address for vidm-secondary-2.

e) For each node, click advanced configuration and click **Select Root Password**.

Select `xint-wsa-root` and click **Save**.

11. On the **Precheck** page, click **Run precheck**.
12. On the **Manual validations** page, select the **I took care of the manual steps above and am ready to proceed** check box and click **Run precheck**.
13. Review the validation report, remediate any errors, and click **Re-run precheck**.
14. Wait for all prechecks to complete with `Passed` messages and click **Next**.
15. On the **Summary** page, review the configuration details. To back up the deployment configuration, click **Export configuration**.
16. To start the deployment, click **Submit**.
The **Request details** page displays the progress of deployment.

17. Monitor the steps of the deployment graph until all stages become `Completed`.

Configure an Anti-Affinity Rule and a Virtual Machine Group for a Clustered Workspace ONE Access Instance

To protect the nodes in a clustered Workspace ONE Access instance from a host-level failure, configure an anti-affinity rule to run the virtual machines on different hosts in the default management vSphere cluster. You then configure a VM group to define the startup order to ensure that vSphere High Availability powers on the clustered Workspace ONE Access nodes in the correct order.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. In the Hosts and Clusters inventory, expand the management domain vCenter Server and data center.
3. Select the cluster and click the **Configure** tab.
4. Create the anti-affinity rule for the clustered Workspace ONE Access virtual machines.
 - a) Navigate to **Configuration > VM/Host rules** and click **Add**.
 - b) Configure the settings and click **OK**.

Setting	Value
Name	<management-domain-name>-anti-affinity-rule-wsa
Enable rule	Selected
Type	Separate Virtual Machines
Members	Click Add , select the clustered Workspace ONE Access nodes, and click OK . <ul style="list-style-type: none"> • vidm-primary_VM • vidm-secondary-1_VM • vidm-secondary-2_VM

5. Create a virtual machine group for the clustered Workspace ONE Access nodes.
 - a) Navigate to **Configuration > VM/Host groups** and click **Add**.
 - b) Configure the settings and click **OK**.

Setting	Value
Name	Clustered Workspace ONE Access Appliances
Type	VM Group
Members	Click Add , select the clustered Workspace ONE Access nodes, and click OK . <ul style="list-style-type: none"> • vidm-primary_VM • vidm-secondary-1_VM • vidm-secondary-2_VM

Configure NTP on Workspace ONE Access

To keep NTP synchronized with the other SDDC components, configure NTP using the Workspace ONE Access appliance configuration interface.

1. In a web browser, log in to the Workspace ONE Access instance with the **admin** user by using the appliance configuration interface (https://<wsa_node_fqdn>:8443/cfg/login).
2. In the navigator pane, click **Time synchronization**.
3. Configure the settings and click **Save**.

Setting	Description
Time sync	NTP selected
NTP Server	Enter the FQDN of the NTP server.

4. If you deployed a cluster, repeat this procedure for the remaining clustered Workspace ONE Access nodes.

Configure the Domain and Domain Search Parameters on Workspace ONE Access

To enable name translation and resolution between the region-specific and the cross-region environments, configure the domain name and domain search parameters on Workspace ONE Access.

1. Log in to the cross-region Workspace ONE Access instance by using a Secure Shell (SSH) client.
2. Switch to the super user by running the **su** command.
3. Open the `/etc/resolv.conf` file in a text editor.

```
vi /etc/resolv.conf
```

4. Add entries for `Domain` and `search` to the end of the file and save the file. For example:

```
Domain rainpole.io
search rainpole.io sfo.rainpole.io
```

5. If you deployed a clustered Workspace ONE Access instance, repeat this procedure for the remaining nodes in the cluster.

Configure an Identity Source for Workspace ONE Access

To enable identity and access management in the SDDC, you integrate your Active Directory with Workspace ONE Access and configure attributes to synchronize users and groups.

1. In a web browser, log in to Workspace ONE Access by using the administration interface to the **System Domain** with **configadmin** user (https://<wsa_fqdn>/admin).
2. On the main navigation bar, click **Identity and access management**.
3. Click the **Directories** tab, and from the **Add directory** drop-down menu, select **Add Active Directory over LDAP/IWA**.
4. On the **Add directory** page, configure the following settings, click **Test connection** and click **Save and next**.

Setting	Value
Directory name	Enter a name for directory. For example, <code>sfo.rainpole.io</code> .
Active Directory over LDAP	Selected
Sync connector	Select the FQDN of

Table continued on next page

Continued from previous page

Setting	Value
	vidm-primary
Do you want this connector to also perform authentication?	Yes
Directory search attribute	SAMAccountName
This Directory requires all connections to use STARTTLS (Optional)	If you want to secure communication between Workspace ONE Access and Active Directory select this option and paste the Root CA certificate in the SSL Certificate box.
Base DN	Enter the Base Distinguished Name from which to start user searches. For example, <code>cn=Users,dc=sfo,dc=rainpole,dc=io.</code>
Bind DN	Enter the DN for the user to connect to Active Directory. For example, <code>cn=svc-wsa-ad,ou=Service Accounts,dc=sfo,dc=rainpole,dc=io.</code>
Bind user password	Enter the password for the Bind user. For example: <code>svc-wsa-ad_password.</code>

- On the **Select the domains** page, review the domain name and click **Next**.
- On the **Map user attributes** page, review the attribute mappings and click **Next**.
- On the **Select the groups (users) you want to sync** page, enter the distinguished name for the folder containing your groups (For example `OU=Security Groups,DC=sfo,DC=rainpole,DC=io`) and click **Select**.
- For each **Group DN** you want to include, select the group to use by Workspace ONE Access for each of the roles, and click **Save** then **Next**.

Product	Role Assigned via Group
Workspace ONE Access	Super Admin
	Directory Admin
	ReadOnly Admin
VMware Aria Suite Lifecycle	VCF Role
	Content Admin
	Content Developers

- On the **Select the Users you would like to sync** page, enter the distinguished name for the folder containing your users (e.g. `OU=Users,DC=sfo,DC=rainpole,DC=io`) and click **Next**.
- On the **Review** page, click **Edit**, from the **Sync frequency** drop-down menu, select **Every 15 minutes**, and click **Save**.
- To initialize the directory import, click **Sync directory**.

Add the Clustered Workspace ONE Access Cluster Nodes as Identity Provider Connectors

To provide high availability for the identity and access management services of a clustered Workspace ONE Access instance, you add the cluster nodes as directory connectors.

This procedure is only applicable if you deployed a clustered Workspace ONE Access instance. It does not apply to a standard Workspace ONE Access instance.

1. In a web browser, log in to the clustered Workspace ONE Access instance by using the administration interface to the **System Domain** with **configadmin** user (https://<wsa_cluster_fqdn>/admin).
2. On the main navigation bar, click **Identity and access management**.
3. Click the **Identity Providers** tab.
4. Click the **WorkspaceIDP__1** identity provider.
5. On the **WorkspaceIDP__1 details** page, under **Connector(s)** from the **Add a connector** drop-down menu, select **vidm-secondary-1_VM**, configure the settings, and click **Add connector**.

Setting	Value
Connector	vidm-secondary-1_VM
Bind to AD	Checked
Bind user password	<i>svc-wsa-ad_password</i>

6. Repeat this step for the **vidm-secondary-2_VM** connector.
7. In the **IdP Hostname** text box, enter the FQDN of the NSX load balancer virtual server for Workspace ONE Access cluster.
8. Click **Save**.

Assign Roles to Active Directory Groups for Workspace ONE Access

Workspace ONE Access uses role-based access control to manage delegation of roles. You assign the **Super Admin**, **Directory Admin** and **ReadOnly** roles to Active Directory groups to manage access to Workspace ONE Access.

You assign the following administrator roles to the corresponding user groups.

Workspace ONE Access Role	Example Active Directory Group Name
Super Admin	wsa-admins
Directory Admin	wsa-directory-admin
ReadOnly Admin	wsa-read-only

1. In a web browser, log in to Workspace ONE Access by using the administration interface to the System Domain with **configadmin** user (https://<wsa_fqdn>/admin).
2. On the main navigation bar, click **Roles**.
3. Assign Workspace ONE Access roles to Active Directory groups.
 - a) Select the **Super Admin** role and click **Assign**.
 - b) In the **Users / User Groups** search box, enter the name of the Active Directory group you want to assign the role to, select the group, and click **Save**.
 - c) Repeat this step to configure the **Directory Admin** and the **ReadOnly Admin** roles.

Assign Roles to Active Directory Groups for VMware Aria Suite Lifecycle

To enable identity and access management for VMware Aria Suite Lifecycle, you integrate the component with the clustered Workspace ONE Access instance.

You assign the following administrative roles to corresponding Active Directory groups.

VMware Aria Suite Lifecycle Role	Example Active Directory Group Name
VCF Role	vrslcm-admins
Content Release Manager	vrslcm-release-manager
Content Developer	vrslcm-content-developer

1. In a web browser, log in to VMware Aria Suite Lifecycle with the **vcfadmin@local** user by using the user interface (https://<vrslcm_fqdn>).
2. On the **My Services** page, click **Identity and Tenant Management**.
3. In the navigation pane, click **User management** and click **Add user / group**.
4. On the **Select users / groups** page, in the search box, enter the name of the group you want to assign the role to, select the Active Directory group, and click **Next**.
5. On the **Select roles** page, select the **VCF Role** role, and click **Next**.
6. On the **Summary** page, click **Submit**.
7. Repeat this procedure to assign roles to the **Content Release Manager** and **Content Developer** user groups.

Working with NSX Federation in VMware Cloud Foundation

With NSX Federation, you can federate NSX environments across VMware Cloud Foundation (VCF) instances. You can manage federated NSX environments with a single pane of glass, create gateways and segments that span VMware Cloud Foundation instances, and configure and enforce firewall rules consistently across instances.

IMPORTANT

If you plan to deploy VMware Aria Suite components, you must deploy Application Virtual Networks before you configure NSX Federation. See [Deploying Application Virtual Networks in VMware Cloud Foundation](#).

NSX Federation is supported between VCF and non-VCF deployments. If you choose to federate NSX between VCF and non-VCF deployments, you are responsible for the deployment and lifecycle of the NSX Global Managers, as well as maintaining version interoperability between VCF-owned NSX Local Managers, non-VCF NSX Local Managers, and the NSX Global Manager.

NSX Federation Key Concepts

NSX Federation introduces some new terms and concepts in VMware Cloud Foundation (VCF).

NSX Federation Systems: Global Manager and Local Manager

An NSX Federation environment within VMware Cloud Foundation includes two types of management systems.

Global Manager: a system similar to NSX Manager that federates multiple Local Managers.

Local Manager: an NSX Manager system in charge of network and security services for a VMware Cloud Foundation instance.

NSX Federation Span: Local and Cross-Instance

When you create a networking object from Global Manager, it can span one or more VMware Cloud Foundation instances.

Local: the object spans only one instance.

Cross-instance: the object spans more than one instance. You do not directly configure the span of a segment. A segment has the same span as the gateway it is attached to.

NSX Federation Tunnel Endpoints

In an NSX Federation environment, there are two types of tunnel endpoints.

Tunnel End Point (TEP): the IP address of a transport node (Edge node or Host) used for Geneve encapsulation within an instance.

Remote Tunnel End Points (RTEP): the IP address of a transport node (Edge node only) used for Geneve encapsulation across instances.

NSX Federation Tier Gateways

An NSX Federation in VMware Cloud Foundation environment includes three types of tier-1 gateways.

Type	Description	Managed By	Scope
standalone tier-1 gateway	Configured in the Local Manager and used for services such as the Load Balancer.	Local Manager	Single VMware Cloud Foundation instance
local-instance tier-1 gateway	Configured in the Global Manager at a single location, this is a global tier-1 gateway used for segments that exist within a single VMware Cloud Foundation Instance.	Global Manager	Single VMware Cloud Foundation instance
cross-instance tier-1 gateway	Configured in the Global Manager, this is a global Tier-1 gateway used for segments that exist across multiple VMware Cloud instances.	Global Manager	Multiple VMware Cloud Foundation instance

Configuring NSX Federation in VMware Cloud Foundation

With NSX Federation, you can federate the management domain NSX or a VI workload domain NSX across VMware Cloud Foundation (VCF) instances.

See [VMware Configuration Maximums](#) for your version of NSX for information about the maximum number of supported federated NSX Managers and other NSX federation maximums.

NOTE

VI workload domains that share an NSX Manager are considered a single location.

Some tasks described in this section are to be performed on the first NSX instance while others need to be performed on each NSX instance that is being federated. See the table below for more information.

NSX Instance	Tasks to be Performed
First Instance	<ol style="list-style-type: none"> 1. Create Global Manager Clusters for VMware Cloud Foundation 2. Replacing Global Manager Cluster Certificates in

Table continued on next page

Continued from previous page

NSX Instance	Tasks to be Performed
	<p>You can skip this step if you are using self-signed certificates.</p> <ol style="list-style-type: none"> 3. Prepare Local Manager for NSX Federation in 4. Enable NSX Federation in 5. Prepare for Stretching Segments between VMware Cloud Foundation Instances: <ol style="list-style-type: none"> a. Create and Configure Cross-Instance Tier-1 Gateway b. Connect Cross-Instance Segments to Cross-Instance Tier-1 Gateway
Enable high availability for NSX Federation Control Plane on one additional instance	<ol style="list-style-type: none"> 1. Create Global Manager Clusters for VMware Cloud Foundation 2. Replacing Global Manager Cluster Certificates in <p>You can skip this step if you are using self-signed certificates.</p> <ol style="list-style-type: none"> 3. Set Standby Global Manager
Each additional instance	<ol style="list-style-type: none"> 1. Prepare Local Manager for NSX Federation in 2. Add Location to Global Manager 3. Stretching Segments between VMware Cloud Foundation Instances: <ol style="list-style-type: none"> a. Delete Existing Tier-0 Gateways in Additional Instances b. Connect Additional VMware Cloud Foundation Instances to Cross-Instance Tier-0 Gateway c. Connect Local Tier-1 Gateway to Cross-Instance Tier-0 Gateway d. Add Additional Instance as Locations to the Cross-Instance Tier-1 Gateway

Create Global Manager Clusters for VMware Cloud Foundation

An NSX Federation environment contains an active and a standby Global Manager cluster and one or more Local Manager clusters. The standby Global Manager appliance provides high availability and disaster recovery.

- [Set Active Global Manager](#)
- [Set Standby Global Manager](#)

Deploy Global Manager Nodes

You deploy three Global Manager nodes in the VMware Cloud Foundation management domain.

1. Download the NSX OVF file from the VMware download portal.
2. In a web browser, log in to vCenter Server at `https://vcenter_server_fqdn/ui`.
3. Select the default cluster in the management domain.
4. Right-click and select **Deploy OVF template**.

5. Select **Local file**, click **Upload files**, and navigate to the OVA file.
6. Click **Next**.
7. Enter a name and a location for the NSX Manager VM, and click **Next**.
The name you enter appears in the vSphere and vCenter Server inventory.
8. Select the compute resource on which to deploy the NSX Manager appliance page and click **Next**.
9. Review and verify the OVF template details and click **Next**.
10. Accept the license agreement and click **Next**.
11. Specify the deployment configuration size and click **Next**.
The Description panel on the right side of the wizard shows the details of selected configuration. You can also refer to [VMware Configuration Maximums](#) to ensure that you choose the correct size for the scale or your environment.
12. Specify storage for the configuration and disk files.
 - Select the virtual disk format.
 - Select the VM storage policy.
 - Specify the datastore to store the NSX Manager appliance files.
 - Click **Next**.

NOTE
The virtual disk format is determined by the selected VM storage policy when using a vSAN datastore.
13. Select the management network as the destination network and click **Next**.
The following steps are all located in the Customize Template section of the Deploy OVF Template wizard.
14. In the Application section, enter the system root, CLI admin, and audit passwords for the NSX Manager. The root and admin credentials are mandatory fields.
Your passwords must comply with the password strength restrictions.
 - At least 12 characters
 - At least one lower-case letter
 - At least one upper-case letter
 - At least one digit
 - At least one special character
 - At least five different characters
15. In the Optional parameters section, leave the password fields blank.
16. In the Network Properties section, enter the hostname of the NSX Manager.
17. For Rolename, select the NSX Global Manager role.
18. Enter the default gateway, management network IPv4, and management network netmask.
19. In the DNS section, enter the DNS Server list and Domain Search list.
20. In the Services Configuration section, enter the NTP Server list and enable SSH.
21. Verify that all your custom OVF template specification is accurate and click **Finish** to initiate the deployment.
The deployment might take 7-8 minutes.
22. After the deployment is complete, power on the Global Manager node.
Right-click the Global Manager VM and, from the **Actions** menu, select **Power > Power on**.
23. In a web browser, log in to Global Manager at `https://gm_node1_fqdn/`.
24. Accept the end-user license agreement and click **Continue**.
25. Join the VMware Customer Experience Program and click **Save**.
26. Repeat steps 4 - 22 to deploy two additional Global Manager nodes.

Join Global Manager Nodes to Form a Cluster

Join the three Global Manager nodes you deployed in the VMware Cloud Foundation management domain to form a cluster.

1. SSH into the first NSX Global Manager node using the `admin` user account.

2. Run the following command to retrieve the Global Manager cluster ID.

```
get cluster config | find Id:
```

3. Copy the output of the command and save it.

4. Run the following command to retrieve the thumbprint of the Global Manager API certificate.

```
get certificate api thumbprint
```

5. Copy the output of the command and save it.

6. Log in to the second Global Manager node and run the following command to join this node to the cluster:

```
join first_node_IP cluster-id cluster_ID username admin password nsx_admin_password  
thumbprint api_thumbprint
```

where `cluster_ID` is the value from step 3 and `certificate_thumbprint` is the value from step 5.

A warning message displays: Data on this node will be lost. Are you sure? (yes/no).

7. Enter `yes` to confirm.

The joining and cluster stabilizing process might take from 10 to 15 minutes.

8. Run `get cluster status` to view the status.

Verify that the status for every cluster service group is UP before making any other cluster changes.

9. Repeat steps 6-8 to join the third node to the cluster.

10. Verify the cluster status on the web interface.

a) Log in to the Global Manager web interface and select **Configuration > Global Manager Appliances**.

b) Verify that the **Cluster status** is green that the cluster node is **Available**.

Create Anti-Affinity Rule for Global Manager Cluster in VMware Cloud Foundation

Create an anti-affinity rule to ensure that the Global Manager nodes run on different ESXi hosts. If an ESXi host is unavailable, the Global Manager nodes on the other hosts continue to provide support for the NSX management and control planes.

1. In a web browser, log in to the management domain or VI workload domain vCenter Server at `https://vcenter_server_fqdn/ui`.

2. Select **Menu > Hosts and Clusters**.

3. In the inventory, expand **vCenter Server > Datacenter**.

4. Select the Global Manager cluster and click the **Configure** tab.

5. Select **VM/Host rules** and click **Add**.

6. Enter the rule details.

Option	Description
Name	Type a name for the rule.

Table continued on next page

Continued from previous page

Option	Description
Enable rule	Select this option.
Type	Select Separate Virtual Machines .
Members	Click Add , select the three Global Manager nodes, and click OK .

- Click **OK** in the Create VM/Host rule dialog box.

Assign a Virtual IP Address to Global Manager Cluster

To provide fault tolerance and high availability to Global Manager nodes, assign a virtual IP address (VIP) to the Global Manager cluster in VMware Cloud Foundation.

- In a web browser, log in to a Global Manager node at https://gm_node_1-fqdn/.
- Click **System** and then select **Global Manager Appliances**.
- Click **Set Virtual IP** and enter the VIP address for the cluster. Ensure that VIP is part of the same subnet as the other management nodes.
- Click **Save**.
- Verify that the VIP is working correctly.

From a browser, log in to the Global Manager using the virtual IP address assigned to the cluster at https://gm_vip_fqdn/.

Prepare Local Manager for NSX Federation in VMware Cloud Foundation

To prepare for NSX Federation, you create an IP pool in the Local Manager. The Global Manager assigns IP addresses from this pool to the Edge nodes for remote tunnel end point (RTEP) interfaces. You also set the global fabric MTU to match the end-to-end MTU between instances.

- In a web browser, log in to Local Manager cluster for the management domain or VI workload domain at https://lm_vip_fqdn/.
- On the main navigation bar, click **Networking**.
- Create an IP pool for RTEP in Local Manager
 - In the navigation pane, select **IP Address Pools** and click **Add IP address pool**.
 - Enter a name.
 - Under Subnets, click **Set**.
 - In the Set Subnets dialog box, click **Add subnet › IP Ranges**.
 - Configure the settings and click **Add**.
 - Click **Add** and then click **Apply**.
 - Click **Save**.
- Configure MTU for RTEP.
 - On the main navigation bar, click **System**.
 - Select **Fabric › Settings**.
 - Under **Global Fabric Settings**, Click **Edit** for Remote Tunnel Endpoint.
 - Enter **9000** in MTU and click **Save**.

Enable NSX Federation in VMware Cloud Foundation

To enable NSX Federation in VMware Cloud Foundation, set the Global Manager as active and add the existing NSX Manager in the management domain or VI workload domain as a location to the Global Manager.

Set Active Global Manager

Activate the Global Manager.

1. In a web browser, log in to Global Manager cluster for the management or VI workload domain at `https://gm_vip_fqdn/`.
2. Click **System** and then select **Location Manager**.
3. Click **Make Active** and enter a name for the active Global Manager.
4. Click **Save**.

Add Location to Global Manager

Add the NSX Manager in the management domain or VI workload domain as a location to the Global Manager. This NSX Manager is now referred to as the Local Manager. You then import segments, tier-0 gateways, and tier-1 gateways from the Local Manager to the Global Manager.

1. Obtain the certificate thumbprint of the NSX Local Manager cluster.
 - a) Enable SSH on one of the NSX Manager VMs.
 - b) From the vCenter UI, open the web console of one of the NSX Managers and login to the Admin user.
 - c) Run the command `start service ssh` to enable SSH on the NSX Manager.
 - d) Use a Secure Shell (SSH) client and log in to the same NSX Manager with the Admin user.
 - e) Run the command `get certificate cluster thumbprint` to retrieve the Local Manager cluster VIP thumbprint.


```
sfo-m01-nsxt01c> get certificate cluster thumbprint

b88c4e052fe61309915527511e7f1b25970286a51cf1dd68ea881daba1ed0a9f
```
 - f) Save the thumbprint.
 - g) Run the `stop service ssh` command to deactivate SSH on the NSX Manager.
2. Add NSX Manager as a location to the Global Manager.
 - a) Log in to Global Manager at `https://active_gm_vip_fqdn/`.
 - b) Select **System > Location Manager** and click **Add On-Prem Location**.
 - c) In the Add New Location dialog box, enter the location details.

Option	Description
Location Name	Enter a name for the location.
FQDN/IP	Enter the FQDN or IP address of the NSX Manager cluster VIP. Do not enter an individual NSX Manager FQDN or IP.
Username and Password	Provide the admin user's credentials for the NSX Manager at the location.
SHA-256 Thumbprint	Add the thumbprint you retrieved in step 1.
Check Compatibility	Click Check Compatibility to ensure that the location can be added. This checks that the NSX version is compatible.

- d) Click **Save**
3. Configure networking on the Local Manager nodes.
 - a) On the Location Manager page, in the Locations section, click **Networking** under the location you are adding then click **Configure**.
 - b) On the Configure Edge Nodes for Stretch Networking page, click **Select All**

- c) In the Remote Tunnel Endpoint Configuration pane enter the following details.

Option	Value
Host Switch	nsxDefaultHostSwitch
Teaming Policy Name	Select Use Default .
RTEP VLAN	Enter the VLAN for the host.
IP Pool for all Nodes	Select the IP pool.

- d) Click **Save**.
4. Import the Local Manager configuration to the Global Manager.
- a) Select the Global Manager context from the drop down menu.

NOTE

You may need to refresh your browser or logout and log in to the Global Manager to see the drop down menu.

- b) On the System tab, select the Location Manager pane.
- c) Under **Locations**, click **Import**.

This option may take 15 minutes or longer to appear.

- d) Verify that you have a recent backup and click **Proceed to import**.
- e) In the Preparing for import dialog box, click **Next** and then click **Import**.
Wait for a confirmation that the import is successful.

Local Manager objects imported into the Global Manager are owned by the Global Manager and appear in the Local Manager with a GM icon. You can modify these objects only from the Global Manager.

5. Repeat these steps for each Local Manager cluster.

Stretch Segments between VMware Cloud Foundation Instances

Each NSX Manager instance to be federated has a tier-0 gateway, tier-1 gateway, and two segments created during NSX Edge deployment and Application Virtual Network (AVN) creation. One of these segments is for local instance use and the other is for cross-instance use. Both segments are initially connected to the same tier-1 gateway. When NSX Manager instances are federated, you create an additional tier-1 gateway for cross-instance use and migrate the cross-instance segment from the original tier-1 gateway to the new tier-1 gateway. The new tier-1 gateway has locations for both instances enabled on it. This allows you to manage the ingress and egress routing for cross-instance segments when you move them between VMware Cloud Foundation instances independently of local instance segments whose ingress and egress remain unaffected.

NOTE

Cross-instance segments cannot have overlapping IP addresses/ranges.

Create and Configure Cross-Instance Tier-1 Gateway

You create a new tier-1 gateway in one of the VMware Cloud Foundation instances. You then extend this gateway to the other federated instances.

1. In a web browser, log in to Global Manager for the management or VI workload domain at https://gm_vip_fqdn/.
2. On the main navigation bar, click **Networking**.
3. In the navigation pane, select **Tier-1 gateways**.
4. Specify the gateway details.

Setting	Specified Value
Tier-1 Gateway Name	Enter a name for the new tier-1 gateway.
Linked Tier-0 Gateway	Enter the global tier-0 gateway.
Edges Pool Allocation Size	Select Routing .
Enable Edge Clusters for Services or Custom span	Select Enabled .
Fail Over	Select Non Preemptive .
Enable Standby Relocation	Select Enabled .
Edge Cluster	Select the Edge cluster.
Mode	Select Primary

5. Click **Save**.
6. Click **Yes** to continue the configuration of the tier-1 gateway.
7. Configure route advertisement for the tier-1 gateway.
 - a) Expand the **Route advertisement** section of the tier-1 gateway.
 - b) Enable all available sources, click **Save**, and click **Close editing**.

Connect Cross-Instance Segments to Cross-Instance Tier-1 Gateway

You connect the cross-instance segments in the first instance to the cross-instance tier-1 gateway you created.

1. In a web browser, log in to Global Manager cluster at https://gm_vip_fqdn/.
2. On the NSX Manager main navigation bar, click **Networking**.
3. In the navigation pane, select **Segments**.
4. On the Segments tab, click the vertical eclipses for the *cross-instance_nsx_segment* and click **Edit**.
5. Change the Connected Gateway from *instance_tier1* to *cross-instance_tier1*, click **Save**, and then click **Close editing**.

Delete Existing Tier-0 Gateways in Additional Instances

Since you will use the cross-instance tier-0 gateway for upstream connections, you delete the local tier-0 gateway from each additional VCF instance.

1. In a web browser, log in to Global Manager cluster at https://active_gm_vip_fqdn/.
2. On the NSX Manager main navigation bar, click **Networking**.
3. Disconnect the tier-1 gateway for the NSX Local Manager.
 - a) In the navigation pane, select Tier-1 Gateways.
 - b) On the Tier-1 Gateways tab, click the vertical eclipses for the *additional_instance_tier1_gateway* and click **Edit**.
 - c) Under Linked Tier-0 gateway, click the X to disconnect the *additional_instance_tier0_gateway*, click **Save**, and click **Close editing**.

CAUTION

At this point any segments connected to *additional_instance_tier1_gateway* will be unreachable until you have finished connecting the additional instance to the cross-instance tier-0 infrastructure.

4. In the navigation pane, select Tier-0 Gateways.

5. On the Tier-0 Gateway page, click the vertical eclipses for the *additional_instance_tier0_gateway* and click **Delete**.
6. Click **Delete**.

Connect Additional VMware Cloud Foundation Instances to Cross-Instance Tier-0 Gateway

You turn the standard tier-0 gateway into a cross-instance tier-0 gateway by connecting additional VMware Cloud Foundation instances to it. You configure uplink interfaces, BGP, and route redistribution for the additional instances.

1. In a web browser, log in to Global Manager cluster at https://active_gm_vip_fqdn/.
2. Add the additional instance as a location on the tier-0 gateway.
 - a) On the NSX Manager main navigation bar, click **Networking**.
 - b) In the navigation pane, select **Tier-0 Gateways**.
 - c) On the Tier-0 Gateway page, click the vertical eclipses for the *cross-instance_tier0_gateway* and click **Edit**.
 - d) Click **Add Location** and enter the required information.

Setting	Value
Location	Select the location name of the instance being added.
Edge Cluster	Select the Edge cluster name of the instance being added.

- e) Click **Save**.
3. Set interfaces for the instance on the tier-0 gateway.
 - a) Expand **Interfaces** and click **Set**.
 - b) Click **Add interface**.
 - c) Enter a name for the interface and select the instance location.
 - d) Set the type to **External** and enter the IP address for the interface.
 - e) Select the segment that the interface is connected to and the Edge node corresponding to the instance.
 - f) Set the MTU to 9000.
 - g) Repeat these steps to add three additional interfaces.
4. Configure BGP neighbors.
 - a) Expand BGP and under BGP Neighbors, click **Set**.
You can enable BFD if the network supports it and is configured for BFD.
 - a) Click **Add BGP neighbor**
 - b) Enter the IP address for the neighbor and select the instance location.
 - c) Enter the remote AS and source addresses for the neighbor.
 - d) Click **Timers & Password** and set the **Hold Down Time** to 12 and **Keep Alive Time** to 4.
 - e) Enter the BGP neighbor password, click **Save**, and then click **Close**.
 - f) Repeat these steps to add another BGP neighbor.
5. Configure Route Re-Distribution
 - a) Expand Route Re-Distribution and next to the location you are adding, click **Set**.
 - b) In the Set Route Re-distribution dialog box, click **Add Route-Redistribution**.
 - c) Enter *default* as name and, under Route re-distribution, click **Set**.
 - d) In the Set route redistribution dialog box, select all listed sources and click **Apply**.
 - e) Click **Add** to finish editing the default route redistribution and click **Apply**.
 - f) Click **Save**

6. Click **Close editing**.

Connect Local Tier-1 Gateway to Cross-Instance Tier-0 Gateway

You connect the local tier-1 gateway at each VCF instance to the cross-instance tier-0 gateway.

1. In a web browser, log in to Global Manager cluster at `https://active_gm_vip_fqdn/`.
2. On the NSX Manager main navigation bar, click **Networking**.
3. In the navigation pane, select **Tier-1 gateways**.
4. On the Tier-1 Gateway page, click the vertical ellipses menu for the `this_instance_tier1_gateway` and click **Edit**.
5. Change the Connected Gateway to `cross_instance_tier0_gateway`.
6. In the Location change dialog box, click **Yes**.
7. Under Locations, delete all locations except the location of the instance you are working with.
8. Click **Save** and click **Close Editing**.

Add Additional Instance as Locations to the Cross-Instance Tier-1 Gateway

Add each additional instance as a location on the cross-instance Tier-1 gateway to enable cross-instance workloads.

1. In a web browser, log in to Global Manager cluster at `https://active_gm_vip_fqdn/`.
2. On the NSX Manager main navigation bar, click **Networking**.
3. In the navigation pane, select **Tier-1 Gateways**.
4. On the Tier-1 Gateway page, click the vertical eclipses for the `cross-instance_tier1` gateway and click **Edit**.
5. Click **Add Location** and enter the following values.

Setting	Value
Location	Select the location of this instance
Edge Cluster	Select the NSX Edge cluster of the this instance
Mode	Set to Secondary .

6. Click **Save** and click **Close Editing**.

Set Standby Global Manager

You provide high availability of the active Global Manager by configuring the Global Manager in the additional instance as standby to the active cluster. In case of failure of the cluster in first instance, you can use the cluster in additional instance to provide the networking capabilities.

Create the standby Global Manager cluster. See [Create Global Manager Clusters for VMware Cloud Foundation](#).

1. Obtain the certificate thumbprint of the Standby Global Manager cluster.
 - a) Enable SSH on one of the NSX Manager VMs.
 - b) From the vCenter UI, open the web console of one of the NSX Managers and login to the Admin user.
 - c) Run the command `start service ssh` to enable SSH on the NSX Manager.
 - d) Use a Secure Shell (SSH) client and log in to the same NSX Manager with the Admin user.
 - e) Run the command `get certificate cluster thumbprint` to retrieve the Global Manager cluster thumbprint.

```
sfo-m01-nsxt01c> get certificate cluster thumbprint
b88c4e052fe61309915527511e7f1b25970286a51cf1dd68ea881daba1ed0a9f
```

- f) Save the thumbprint.
 - g) Run the **stop service ssh** command to deactivate SSH on the NSX Manager.
2. Add additional Global Manager instance
 - a) Log in to the Active Global Manager at https://active_gm_vip_fqdn/.
 - b) On the main navigation bar, Select **System > Location Manager**.
 - c) Click **Add Standby**.
 - d) Enter the location name, FQDN, username and password, and the SHA-256 thumbprint you had retrieved earlier.
 - e) Click **Check Compatibility** and click **Save**.

Replacing Global Manager Cluster Certificates in VMware Cloud Foundation

To replace certificates for the Global Manager cluster, you import root and intermediate CA-signed certificates as appropriate and replace the Global Manager default certificates with the imported certificates using API calls.

Import a CA-Signed Certificate to the Global Manager Cluster

Import the root/leaf or machine certificate and intermediate certificate as appropriate to the first Global Manager node.

Generate root and intermediate CA-signed certificates.

1. In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
2. Import the root CA certificate.
 - a) On the main navigation bar, **System > Certificates**.
 - b) Click **Import > Import CA certificate**.
 - c) In the Import CA Certificate dialog box, enter a name for the root CA certificate.
 - d) For **Certificate Contents**, select the root CA certificate you created in step 2c and click **Import**.
3. Import certificates for the Global Manager nodes and the load balanced virtual server address.
 - a) Click **Import > Import certificate**.
 - b) In the **Name** field, enter *gm_vip_fqdn*.
 - c) In the Certificate Contents, browse to the previously created certificate file with the extension `chain.pem` and select the file.
 - d) In the **Private Key**, browse to the previously created private key with the extension `.key`, select the file, and click **Import**.

Replace the Certificate for the First Global Manager Node

Replace the default certificate of the first Global Manager node to establish a trusted connection with the management components in the SDDC. You use APIs for this procedure.

1. In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
2. Retrieve the certificate ID.
 - a) On the main navigation bar, click **System > Certificates**.
 - b) Copy the certificate ID value and save it.
3. Log in to the host that has access to your data center.
4. Replace the default certificate on the first Global Manager node with the CA-signed certificate.
 - a) Start the Postman application in your web browser and log in.
 - b) On the **Authorization** tab, enter the following settings.

Setting	Value
Type	Select Basic Auth .
User name	Enter <code>admin</code> .
Password	Enter <code>nsx_admin_password</code> .

- c) Click **Update request**.
d) On the Headers tab, add a key as follows.

Setting	Value
Key	Content-Type
Key Value	application/xml

- e) In the request pane at the top, send the following HTTP request.

Setting	Value
HTTP request method	Select POST .
URL	Enter <code>https://gm_node1_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>

After the Global Manager sends a response, a 200 OK status is displayed on the Body tab.

5. Restart the first Global Manager node.
- Log in to vCenter Server.
 - In the inventory expand **vCenter Server** › **Datacenter** › **Cluster**.
 - Right-click the node and select **Actions** › **Power** › **Restart guest OS**.

Replace Certificates and Virtual IP for the Remaining Global Manager Nodes

Replace the default certificates on the remaining Global Manager nodes.

Table 230: URLs for Replacing the Global Manager Node Certificates

NSX Manager Node	POST URL for Certificate Replacement
<code>gm_node2_fqdn</code>	<code>https://gm_node2_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>
<code>gm_node3_fqdn</code>	<code>https://gm_node3_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=gm_fqdn_certificate_ID</code>
<code>gm_vip_fqdn</code>	<code>https://gm_vip_fqdn/api/v1/cluster/api-certificate?action=set_cluster_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>

1. In a web browser, log in to the active Global Manager at `https://gm_vip_fqdn/`.

2. Log in to the host that has access to your data center.
3. Replace the default certificate for the second Global Manager node with the CA-signed certificate by using the first Global Manager node as a source.
 - a) Start the Postman application in your web browser and log in.
 - b) On the **Authorization** tab, configure the following settings.

Setting	Value
Type	Select Basic Auth.
User name	Enter <code>admin</code> .
Password	Enter the <code>nsx_admin_password</code> .

- a) Click **Update request**.
- b) On the **Headers** tab, enter the header details.

Setting	Value to Select
Key	Content-Type
Key Value	application/xml

- c) In the request pane at the top, send the URL query.

Setting	Value
HTTP request method	Select POST.
URL	Enter <code>https://gm_node2_fqdn/api/v1/node/services/http?action=apply_certificate&certificate_id=firstinstance_gm_vip_certificate_ID</code>

After the NSX Manager appliance responds, the Body tab displays a 200 OK status.

4. To upload the CA-signed certificate on the third Global Manager node, repeat steps 2 to step 4 with appropriate values.
5. Restart the second and third Global Manager nodes.
 - a) Log in to vCenter Server.
 - b) In the inventory expand **vCenter Server** › **Datacenter** › **Cluster**
 - c) Right-click the second and third Global Manager nodes and click **Actions** › **Power** › **Restart guest OS**.
6. Verify the status of each Global Manager node.
 - a) In a web browser, log in to the first Global Manager node at `https://gm_node1_fqdn/`.
 - b) For each node, navigate to **System** › **Global Manager Appliances** › **View Details** and confirm that the status is **REPO_SYNC = SUCCESS**.
7. Assign a certificate to the Global Manager cluster.
 - a) Start the Postman application in your web browser and log in.
 - b) On the **Authorization** tab, configure the following settings.

Setting	Value
Type	Select Basic Auth .
User name	Enter <code>admin</code> .
Password	Enter <code>nsx_admin_password</code> .

- c) Click **Update request**.
- d) On the Headers tab, add a key as follows.

Setting	Value
Key	Content-Type
Key Value	application/xml

- e) In the request pane at the top, send the URL query.

Setting	Value
HTTP request method	Select POST .
URL	Enter <code>https://gm_vip_fqdn/api/v1/cluster/api-certificate?action=set_cluster_certificate&certificate_id=gm_vip_fqdn_certificate_ID</code>

After the NSX Global Manager sends a response, a 200 OK status is displayed on the Body tab.

Update Local Manager Certificate Thumbprint in Global Manager Cluster

After you rotate the Local Manager certificates using SDDC Manager, you obtain the new certificate thumbprint to update it in the Global Manager cluster.

1. In a web browser, log in to Global Manager at `https://nsx_gm_vip_fqdn/`.
2. Obtain certificate thumbprint.
 - a) Log in to a vCenter Server by using a Secure Shell (SSH) client.
 - b) Run the **shell** command to switch to the bash shell.
 - c) Run the command to retrieve the SHA-256 thumbprint of the virtual IP for the NSX Manager cluster certificate.

```
echo -n | openssl s_client -connect nsx_lm_vip_fqdn:443 2>/dev/null | openssl x509 -noout -fingerprint -sha256
```

- d) Save the thumbprint value.
3. Update the Local Manager certificate thumbprint in the Global Manager.
 - a) On the main navigation bar, click **System**.
 - b) In the navigation pane, select **Location Manager**.
 - c) Under **Locations**, select the Local Manager instance, and click **Actions**.
 - d) Click **Edit Settings** and update NSX Local Manager Certificate Thumbprint.
 - e) Click **Check Compatibility** and click **Save**.
 - f) Wait for the Sync Status to display success and verify that all Local Manager nodes appear.

- Under Locations, update the Local Manager certificate thumbprint for all the instances.

Password Management for NSX Global Manager Cluster in VMware Cloud Foundation

You can manage NSX Global Manager user accounts using the NSX appliance's CLI. Resetting the password for any of the local users on one node automatically resets the password for the other NSX Managers in the cluster. The synchronization of the password can take a few minutes.

See [Manage Local User's Password or Name Using the CLI](#).

Backup and Restore of NSX Global Manager Cluster in VMware Cloud Foundation

Regular backups of the NSX Global Manager components ensures that you can keep your environment operational if a data loss or failure occurs.

The Global Manager cluster stores the configured state of the segments. If the Global Manager appliances become unavailable, the network traffic in the data plane is intact but you can make no configuration changes.

Configure NSX Global Manager Cluster Backups

Configure an SFTP server to store backup files. After a backup file server is configured, you can start a backup at any time, or schedule recurring backups.

- In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
- Select **System > Backup & Restore**.
- On the Backup tab, click **Edit**.
- Enter the IP address or FQDN of the backup file server.
- Change the default port if necessary. The default port is 22.
- The protocol text box is already filled in. SFTP is the only supported protocol.
- In the **Directory Path** text box, enter the absolute directory path where the backups will be stored.
- Enter the user name and password required to log in to the backup file server.

The first time you configure a file server, you must provide a password. Subsequently, if you reconfigure the file server, and the server IP or FQDN, port, and user name are the same, you do not need to enter the password again.

- Leave the **SSH Fingerprint** blank and accept the fingerprint provided by the server after you click Save in a later step.
- Enter a passphrase.

NOTE

You will need this passphrase to restore a backup. If you forget the passphrase, you cannot restore any backups.

- Click Edit under the Schedule label.

You can schedule recurring backups or trigger backups for configuration changes.

- Click the Recurring Backup toggle.
- Click Weekly and set the days and time of the backup, or click Interval and set the interval between backups.
- Enabling the **Detect NSX configuration change** option will trigger an unscheduled full configuration backup when it detects any runtime or non-configuration related changes, or any change in user configuration. For Global Manager, this setting triggers backup if any changes in the database are detected, such as the addition or removal of a Local Manager or Tier-0 gateway or DFW policy.

4. You can specify a time interval for detecting database configuration changes. The valid range is 5 minutes to 1,440 minutes (24 hours). This option can potentially generate a large number of backups. Use it with caution.
5. Click **Save**.

After you configure a backup file server, you can click **Backup Now** to manually start a backup at any time. Automatic backups run as scheduled. You see a progress bar of your in-progress backup.

Restore an NSX Global Manager Cluster Backup

Restoring a backup restores the state of the network at the time of the backup. In addition, the configurations maintained by Global Manager appliances are also restored.

- Verify that you have the login credentials for the backup file server.
- Verify that you have the SSH fingerprint of the backup file server. Only SHA256 hashed ECDSA (256 bit) host key is accepted as a fingerprint.
- Verify that you have the passphrase of the backup file.

Do not change the configuration of the NSX Global Manager cluster while the restore process is in progress.

1. If any nodes in the appliance cluster that you are restoring are online, power them off.
2. Install one new appliance node on which to restore the backup.
 - If the backup listing for the backup you are restoring contains an IP address, you must deploy the new Global Manager node with the same IP address. Do not configure the node to publish its FQDN.
 - If the backup listing for the backup you are restoring contains an FQDN, you must configure the new appliance node with this FQDN and publish the FQDN. Only lowercase FQDN is supported for backup and restore.
3. In a web browser, log in to Global Manager at https://gm_vip_fqdn/.
4. Make the Global Manager active. You can restore a backup only on an active Global Manager.
 - a) On the main navigation bar, click **System**.
 - b) In the navigation pane, select **Location Manager**.
 - c) On the Location Manager page, click **Make Active**, enter a name for the Global Manager, and click **Save**.
5. On the main navigation bar, click **System > Backup & Restore** and then click **Edit**.
6. Enter the IP address or FQDN of the backup file server.
7. Change the default port if necessary. The default port is 22.
8. To log in to the server, enter the user name and password.
9. In the **Destination Directory** text box, enter the absolute directory path where the backups are stored.
10. Enter the passphrase that was used to encrypt the backup data.
11. Leave the **SSH Fingerprint** blank and accept the fingerprint provided by the server after you click Save in a later step.
12. Select a backup and click **Restore**.
13. The restore process prompts you to take action, if necessary, as it progresses.
14. After the restored manager node is up and functional, deploy additional nodes to form a NSX Global Manager cluster.

Stretching vSAN Clusters in VMware Cloud Foundation on Dell VxRail

You can stretch a vSAN cluster in a workload domain across two availability zones within a region. Both availability zones must contain an equal number of hosts to ensure failover in case any of the availability zones goes down.

The default management cluster must be stretched before a VI workload domain cluster can be stretched. This ensures that the NSX control plane and management VMs (vCenter, NSX, SDDC Manager) remain accessible if the stretched cluster in the second availability zone goes down.

NOTE

You cannot stretch a cluster in the following conditions:

- The cluster uses vSAN ESA.

NOTE

Starting with VMware Cloud Foundation 5.2.1.1 you can stretch a cluster that uses vSAN ESA. Earlier versions of VMware Cloud Foundation only support stretching vSAN OSA clusters.

- The cluster has a vSAN remote datastore mounted on it.
- The cluster shares a vSAN Storage Policy with any other clusters.

You may want to stretch a cluster for the following reasons.

- Planned maintenance

You can perform a planned maintenance on an availability zone without any downtime and then migrate the applications after the maintenance is completed.

- Automated recovery

Stretching a cluster automatically initiates VM restart and recovery, and has a low recovery time for the majority of unplanned failures.

- Disaster avoidance

With a stretched cluster, you can prevent service outages before an impending disaster.

This release of VMware Cloud Foundation does not support deleting or unstretching a cluster.

About Availability Zones and Regions

This section describes availability zones and regions as used for stretch clusters.

Availability Zones

An availability zone is a collection of infrastructure components. Each availability zone runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable. Each zone should have independent power, cooling, network, and security.

Additionally, these zones should be physically separate so that disasters affect only one zone. The physical distance between availability zones is short enough to offer low, single-digit latency (less than 5 ms) and large bandwidth (10 Gbps) between the zones.

Availability zones can either be two distinct data centers in a metro distance, or two safety or fire sectors (data halls) in the same large-scale data center.

Regions

Regions are in two distinct locations - for example, region A can be in San Francisco and region B in Los Angeles (LAX). The distance between regions can be rather large. The latency between regions must be less than 150 ms.

Stretched Cluster Requirements

In an environment with multiple availability zones, Layer 2 networks must be stretched between the availability zones by the physical infrastructure. You also must provide a Layer 3 gateway that is highly available between availability zones. The method for stretching these Layer 2 networks and providing a highly available Layer 3 gateway is vendor-specific.

VLANs and Subnets for Multiple Available Zones

This section displays a sample configuration for an environment with multiple availability zones. The VM management, Uplink 01, Uplink 02, and Edge overlay networks in each availability zone must be stretched to facilitate failover of the

NSX Edge appliances between availability zones. The Layer 3 gateway for the management and Edge overlay networks must be highly available across the availability zones.

NOTE

If VLAN is stretched between AZ1 and AZ2, the Layer 3 network must also be stretched between the two AZs.

Table 231: Stretched Cluster Subnet Requirements

Function	Availability Zone 1	Availability Zone 2	HA Layer 3 Gateway	Recommended MTU
VM Management VLAN	✓	✓	✓	1500
Management VLAN (AZ1)	✓	X	✓	1500
vMotion VLAN	✓	X	✓	9000
vSAN VLAN (AZ1)	✓	X	✓	9000
NSX Host Overlay VLAN	✓	X	✓	9000
NSX Edge Uplink01 VLAN	✓	✓	X	9000
NSX Edge Uplink02 VLAN	✓	✓	X	9000
NSX Edge Overlay VLAN	✓	✓	✓	9000
Management VLAN (AZ2)	X	✓	✓	1500
vMotion VLAN (AZ2)	X	✓	✓	9000
vSAN VLAN (AZ2)	X	✓	✓	9000
NSX Host Overlay VLAN (AZ2)	X	✓	✓	9000

Networking for Multiple Availability Zones

There are specific physical data center network requirements for a topology with multiple availability zones. For information about the vSAN witness appliance requirements, see [vSAN Witness Design for](#) in the *VMware Cloud Foundation Design Guide*.

Table 232: Physical Network Requirements for Multiple Availability Zone

Component	Requirement
MTU	<p>VLANs which are stretched between availability zones must meet the same requirements as the VLANs for intra-zone connection including MTU. MTU value must be consistent end-to-end including components on the inter-zone networking path. Set MTU values as follows.</p> <ul style="list-style-type: none"> • MTU for all VLANs and Switch Virtual Interfaces (vMotion, Geneve, and Storage) to jumbo frames. • Management MTU to 1500. • Geneve overlay requires a minimum MTU of 1600.

Table continued on next page

Continued from previous page

Component	Requirement
Layer 3 gateway availability	For VLANs that are stretched between available zones, configure data center provided method to failover the Layer 3 gateway between availability zones. For example, VRRP or HSRP.
DHCP availability	For VLANs that are stretched between availability zones, provide high availability for the DHCP server so that a failover operation of a single availability zone will not impact DHCP availability.
BGP routing	Each availability zone data center must have its own Autonomous System Number (ASN).
Ingress and egress traffic	<ul style="list-style-type: none"> • For VLANs that are stretched between availability zones, traffic flows in and out of a single zone. Local egress is not supported. • For VLANs that are not stretched between availability zones, traffic flows in and out of the zone where the VLAN is located. • For NSX virtual network segments that are stretched between regions, traffic flows in and out of a single availability zone. Local egress is not supported.
Latency	<p>vSphere</p> <ul style="list-style-type: none"> • Less than 150 ms latency RTT for vCenter Server connectivity. • Less than 150 ms latency RTT for vMotion connectivity. • Less than 5 ms latency RTT for vSAN hosts connectivity. <p>vSAN</p> <ul style="list-style-type: none"> • Less than 200 ms latency RTT for up to 10 hosts per site. • Less than 100 ms latency RTT for 11-15 hosts per site. <p>NSX Managers</p> <ul style="list-style-type: none"> • Less than 10 ms latency RTT between NSX Managers • Less than 150 ms latency RTT between NSX Managers and transport nodes.

Deploy and Configure vSAN Witness Host

Each vSAN stretched cluster requires a witness host deployed in a vSAN witness zone, which must be different from the location of both availability zones.

You deploy the vSAN witness host using an appliance instead of using a dedicated physical ESXi host as a witness host. The witness host does not run virtual machines and must run the same version of ESXi as the ESXi hosts in the stretched cluster. It must also meet latency and Round Trip Time (RTT) requirements.

There are separate vSAN witness appliances for vSAN OSA and vSAN ESA. You must deploy the witness appliance that matches the cluster type that you are stretching.

See the Physical Network Requirements for Multiple Availability Zone table within [Stretched Cluster Requirements](#).

Deploy vSAN Witness Host

You deploy the vSAN witness host for a stretched cluster at a site which is isolated from the existing availability zones to prevent propagation of failure or outage in the data center.

Download the VMware vSAN Witness Appliance .ova file from the [Broadcom Support Portal](#).

For more information, see [vSAN Witness Design for VMware Cloud Foundation](#).

1. In a web browser, log in to vCenter Server at `https://vcenter_server_fqdn/ui`.
2. Select **Menu** › **Hosts and Clusters**.
3. In the inventory panel, expand **vCenter Server** › **Datacenter**.
4. Right-click the cluster and select **Deploy OVF template**.
5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, browse to the location of the vSAN witness host OVA file, and click **Next**.
6. On the **Select a name and folder** page, enter a name for the virtual machine and click **Next**.
7. On the **Select a compute resource** page, click **Next**.
8. On the **Review details** page, review the settings and click **Next**.
9. On the **License agreements** page, accept the license agreement and click **Next**.
10. On the **Configuration** page, select **Medium** and click **Next**.
11. On the **Select storage** page, select a datastore and click **Next**.
12. On the **Select networks** page, select a portgroup for the witness and management network, and click **Next**.
13. On the **Customize template** page, enter the root password for the witness and click **Next**.
14. On the **Ready to complete** page, click **Finish** and wait for the process to complete.
15. Power on the vSAN witness host.
 - a) In the inventory panel, navigate to **vCenter Server** › **Datacenter** › **Cluster**.
 - b) Right-click the vSAN witness host and from the **Actions** menu, select **Power** › **Power on**.

Configure the Management Network on the vSAN Witness Host

Configure the management network for the vSAN witness host in the ESXi Direct Console User Interface (DCUI).

1. In the inventory panel of the vCenter Server Client, select **vCenter Server** › **Datacenter**.
2. Open the DCUI of the ESXi host.
 - a) Right-click the vSAN witness host and click **Open remote console**.
 - b) Press F2 to enter the DCUI.
 - c) Log in with the `vsan_witness_root_password`.
3. Configure the network.
 - a) Select **Configure Management Network** and press Enter.
 - b) Select **IPv4 Configuration** and press Enter.
 - c) Select **Set static IPv4 address and network configuration** and press the Space bar.
 - d) Enter **IPv4 Address**, **Subnet Mask** and **Default Gateway** and press Enter.
 - e) Select **DNS Configuration** and press Enter.
 - f) Select **Use the following DNS Server address and hostname** and press the Space bar.
 - g) Enter **Primary DNS Server**, **Alternate DNS Server** and **Hostname** and press Enter.
 - h) Select **Custom DNS Suffixes** and press Enter.
 - i) Ensure that there are no suffixes listed and press Enter.
4. Press Escape to exit and press Y to confirm the changes.

Register vSAN Witness Host

Before you can configure the vSAN Witness Host, you must register it with vCenter Server.

1. Use the vSphere Client to log in to the vCenter Server containing the cluster that you want to stretch.
2. In the vSphere Client, navigate to the data center.
3. Right-click the data center and select **Add Host**.

IMPORTANT

You must add the vSAN Witness Host to the datacenter. Do not add it to a folder.

4. On the **Name and location** page, enter the Fully Qualified Domain Name (FQDN) of the vSAN Witness Host and click **Next**.

NOTE

Do not use the IP address.

5. On the **Connection settings** page, enter administrator credentials and click **Next**.
6. On the **Host summary** page, review the summary of the host details and click **Next**.
7. On the **Host lifecycle** page, the check box **Manage host with an image** is selected by default.
 - If you want to manage the host with an image, leave the check box selected and click **Next**.
 - If you do not want to manage the host with an image, deselect the check box and click **Next**.
8. If you manage the host with an image, on the **Image** page, set up the desired image and click **Next**.
9. On the **Assign license** page, assign an existing license and click **Next**.

NOTE

Do not create a new license.

10. Review the summary and click **Finish**.

Configure NTP on the Witness Host

To prevent time synchronization issues, configure the NTP service on the vSAN witness host.

1. In the inventory panel of the vCenter Server Client, select **vCenter Server > Datacenter**.
2. Select the vSAN witness host and click the **Configure** tab.
3. Configure the NTP client on the vSAN witness host.
 - a) In the **System** section, click **Time configuration** and click the **Edit** button.
 - b) Select **Use Network Time Protocol (enable NTP client)**.
 - c) Configure the following settings and click **OK**.

Setting	Value
NTP Servers	NTP server address
Start NTP Service	Selected
NTP Service Startup Policy	Start and stop with host

Configure the VMkernel Adapters on the vSAN Witness Host

To enable vSAN data network communication between the availability zones, configure the witness network on the vSAN witness host.

1. In the inventory panel of the vCenter Server Client, select **vCenter Server > Datacenter**.
2. Select the vSAN witness host and click the **Configure** tab.
3. Remove the dedicated witness traffic VMkernel adapter on the vSAN Witness host.
 - a) In the **Networking** section, click **VMkernel adapters**.
 - b) Select the kernel adapter **vmk1** with `secondaryPg` as **Network label** and click **Remove**.
 - c) On the **Remove VMkernel adapter** dialog box, click **Remove**
4. Remove the virtual machine network port group on the vSAN witness host.
 - a) In the left pane, select **Networking > Virtual switches**.
 - b) Expand the **Standard switch: secondary switch** section.
 - c) Click the vertical ellipsis and from the drop-down menu, select **Remove**.
 - d) On the **Remove standard switch** dialog box, click **Yes**.
 - e) Expand the **Standard switch: vSwitch0** section.
 - f) In the **VM Network** pane, click the vertical ellipsis and from the drop-down menu, select **Remove**.
 - g) On the **Remove port group** dialog box, click **Yes**.
5. Enable witness traffic on the VMkernel adapter for the management network of the vSAN witness host.
 - a) On the **VMkernel adapters** page, select the **vmk0** adapter and click **Edit**.
 - b) In the **vmk0 - edit settings** dialog box, click **Port properties**, select the **vSAN** check box, and click **OK**.

Stretch a VxRail Cluster in VMware Cloud Foundation

This procedure describes how to stretch a VxRail cluster across two availability zones. You can stretch a vSAN cluster in the management domain or VI workload domain.

- Verify that vCenter Server is operational.
- Verify that you have completed the Planning and Preparation Workbook with the management domain or VI workload domain deployment option included.
- Verify that your environment meets the requirements listed in the Prerequisite Checklist sheet in the Planning and Preparation Workbook.
- Ensure that you have enough hosts such that there is an equal number of hosts on each availability zone. This is to ensure that there are sufficient resources in case an availability zone goes down completely.
- Deploy and configure a vSAN witness host. See [Deploy and Configure vSAN Witness Host](#).
- If you are stretching a cluster in a VI workload domain, the default management vSphere cluster must have been stretched.
- Download <https://community.broadcom.com/vmware-code/viewdocument/vcf-on-vxrail-stretch-cluster-7>.

NOTE

You cannot stretch a cluster in the following conditions:

- The cluster uses vSAN ESA.

NOTE

Starting with VMware Cloud Foundation 5.2.1.1 you can stretch a cluster that uses vSAN ESA. Earlier versions of VMware Cloud Foundation only support stretching vSAN OSA clusters.

- The cluster has a vSAN remote datastore mounted on it.
- The cluster shares a vSAN Storage Policy with any other clusters.

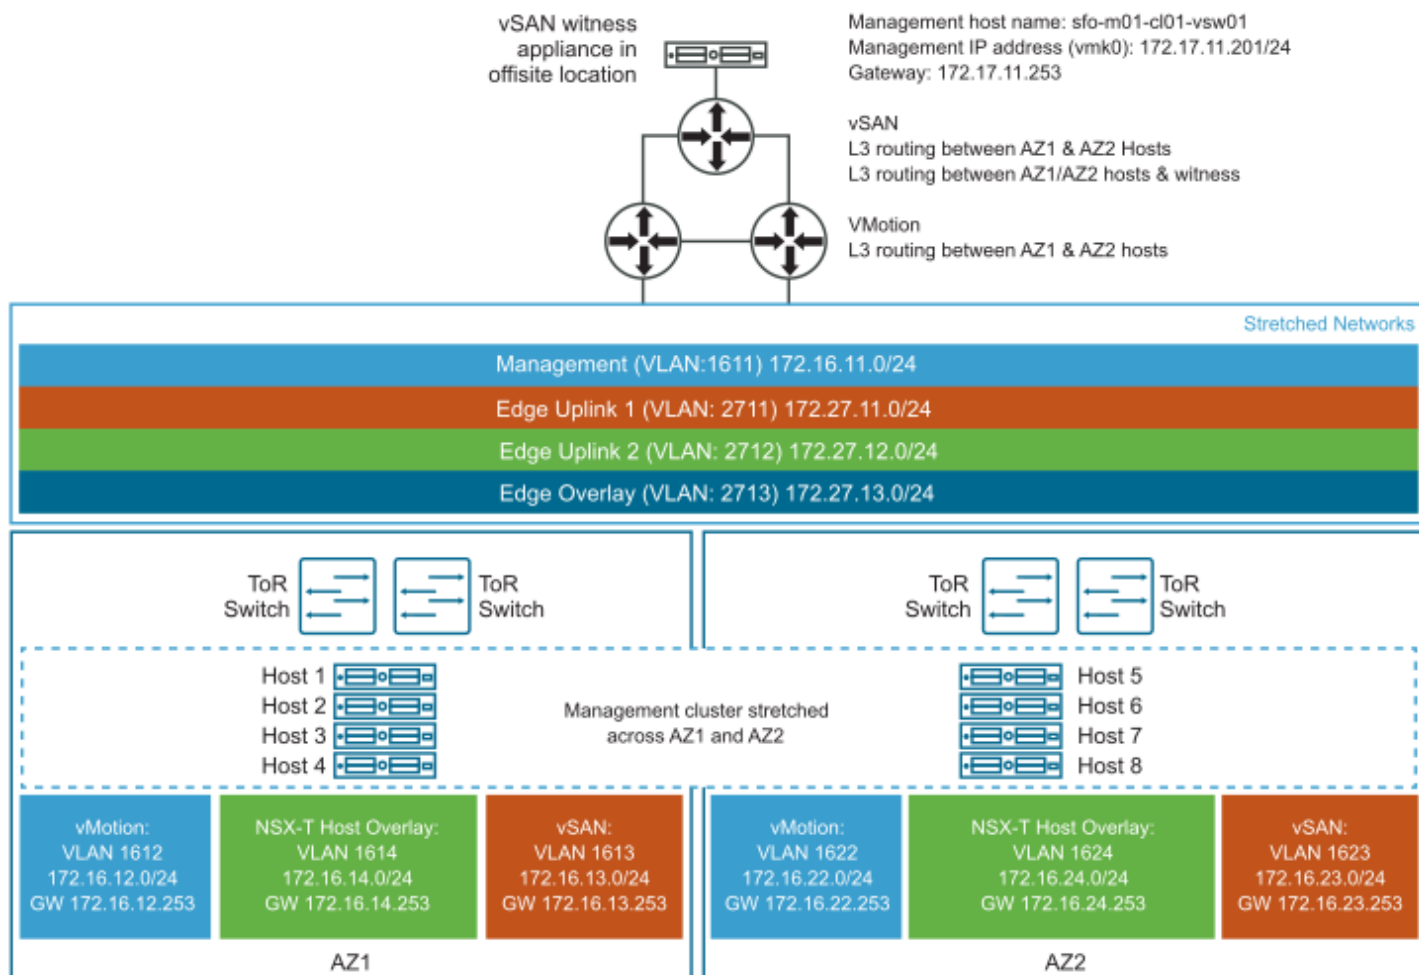
When you stretch a cluster, VMware Cloud Foundation modifies the site disaster tolerance setting for storage policy associated with datastore of that cluster from **None - standard cluster** to **Site mirroring - stretched cluster**. This affects

all VMs using default datastore policy in that cluster. If you do not want to change the site disaster tolerance setting for specific VMs, apply a different storage policy to those VMs before stretching the cluster. This example use case has two availability zones in two buildings in an office campus - AZ1 and AZ2. Each availability zone has its own power supply and network. The management domain is on AZ1 and contains the default cluster, SDDC-Cluster1. This cluster contains four ESXi hosts.

VSAN network	VLAN ID=1623
	MTU=9000
	Network=172.16.234.0
	netmask 255.255.255.0
	gateway 172.16.23.253
	IP range=172.16.23.11 - 172.16.234.59
vMotion network	VLAN ID=1622
	MTU=9000
	Network=172.16.22.0
	netmask 255.255.255.0
	gateway 172.16.22.253
	IP range=172.16.22.11 - 172.16.22.59

There are four ESXi hosts in AZ2 that are not in the VMware Cloud Foundation inventory yet. We will stretch the default cluster SDDC-Cluster1 in the management domain from AZ1 to AZ2.

Figure 24: Stretch Cluster Example



To stretch a cluster for VMware Cloud Foundation on Dell VxRail, perform the following steps:

1. Using an SSH File Transfer tool, copy `initiate_stretch_cluster_vxrail_<version>.py` to the `/home/vcf/` directory on the SDDC Manager appliance.
2. Using SSH, log in to the SDDC Manager appliance with the user name `vcf` and the password you specified in the deployment parameter workbook.
3. Run the script with `-h` option for details about the script options.

```
python initiate_stretch_cluster_vxrail_<version>.py -h
```

4. Run the following command to prepare the cluster to be stretched. The command creates affinity rules for the VMs to run on the preferred site:

```
python initiate_stretch_cluster_vxrail_<version>.py --workflow prepare-stretch --sc-domain<SDDC-valid-domain-name>--sc-cluster<valid-cluster-name>
```

Replace `<SDDC-valid-domain-name>` and `<valid-cluster-name>` with the correct values for your environment. For example:

```
python initiate_stretch_cluster_vxrail_<version>.py --workflow prepare-stretch --sc-domain wdc1-workflowspec-vxrail --sc-cluster VxRail-Virtual-SAN-Cluster-8d2c9f37-e230-4238-ab35-cafd5033a59e
```

Enter the SSO user name and password when prompted to do so.

Once the workflow is triggered, track the task status in the SDDC Manager UI. If the task fails, debug and fix the issue and retry the task from the SDDC Manager UI. Do not run the script again.

5. Use the VxRail vCenter plug-in to add the additional hosts in Availability Zone 2 to the cluster by performing the VxRail Manager cluster expansion work flow.
6. Run the following command to stretch the cluster:

```
python initiate_stretch_cluster_vxrail_<version>.py --workflow stretch-vsan --sc-domain<SDDC-valid-domain-name>--sc-cluster<valid cluster name which is a part of the domain to be stretched>--sc-hosts<valid host names>--witness-host-fqdn<witness host/appliance IP or fqdn>--witness-vsan-ip<witness vsan IP address>--witness-vsan-cidr<witness-vsan-network-IP-address-with-mask>
```

Replace *<SDDC-valid-domain-name>*, *<valid cluster name which is a part of the domain to be stretched>*, *<valid host names>*, *<witness vsan IP address>*, *<witness host/appliance IP or fqdn>*, *<witness vsan IP address>*, and *<witness-vsan-network-IP-address-with-mask>* with the correct values for your environment. For example:

```
python initiate_stretch_cluster_vxrail_<version>.py --workflow stretch-vsan --sc-domain wdc1-workflowspec-vxrail --sc-cluster VxRail-Virtual-SAN-Cluster-8d2c9f37-e230-4238-ab35-cafd5033a59e --sc-hosts wdc3-005-proxy.vxrail.local --witness-host-fqdn 172.16.10.235 --witness-vsan-ip 172.16.20.235 --witness-vsan-cidr 172.16.20.0/24
```

7. When prompted, enter the following information:
 - SSO user name and password
 - Root user password for ESXi hosts
 - vSAN gateway IP for the preferred (primary) and non-preferred (secondary) site
 - vSAN CIDR for the preferred (primary) and non-preferred (secondary) site
 - VLAN ID for the non-preferred site overlay VLAN
 - IP address pool details (if supported and required)
 - Confirm the SSH thumbprints for the hosts

Once the workflow is triggered, the task is tracked in the SDDC Manager UI. If the task fails, debug and fix the issue and retry from SDDC Manager UI. Do not run the script again.

8. Monitor the progress of the AZ2 hosts being added to the cluster.
 - a) In the SDDC Manager UI, click **View All Tasks**.
 - b) Refresh the window to monitor the status.
9. Validate that stretched cluster operations are working correctly by logging in to the vSphere Web Client.
 - a) Verify vSAN Health.
 1. On the home page, click **Host and Clusters** and then select the stretched cluster.
 2. Click **Monitor** > **vSAN** > **Skyline Health**.
 3. Click **Retest**.
 4. Fix errors, if any.

- b) Verify the vSAN Storage Policy.
1. On the home page, click **Policies and Profiles** › **VM Storage Policies** › **vSAN Default Storage Policies**.
 2. Select the policy associated with the vCenter Server for the stretched cluster and click **Check Compliance**.
 3. Click **VM Compliance** and check the **Compliance Status** column for each VM.
 4. Fix errors, if any.

NSX Configuration for Availability Zone 2

To provide the necessary networking services for fail-over of SDDC components from availability zone 1 to availability zone 2 in the management domain, you configure NSX for availability zone 2.

Configure IP Prefixes in the Tier-0 Gateway for Availability Zone 2

You configure default and any IP prefixes on the tier-0 gateway to permit access to route advertisement by any network and by the 0.0.0.0/0 network. These IP prefixes are used in route maps to prepend a path to one or more autonomous systems (AS-path prepend) for BGP neighbors and to configure local-reference on the learned default-route for BGP neighbors in availability zone 2.

1. In a web browser, log in to NSX Manager for the management or workload domain to be stretched at `https://nsx_manager_fqdn/login.jsp?local=true`.
2. On the main navigation bar, click **Networking**.
3. In the navigation pane, click **Tier-0 gateways**.
4. Select the gateway and from the ellipsis menu, click **Edit**.
5. Create the Any IP prefix list.
 - a) Expand the **Routing** section and in the **IP prefix list** section, click **Set**.
 - b) In the **Set IP prefix list** dialog box, click **Add IP prefix list**.
 - c) Enter `Any` as the prefix name and under **Prefixes**, click **Set**.
 - d) In the **Set prefixes** dialog box, click **Add Prefix** and configure the following settings.

Setting	Value
Network	any
Action	Permit

- e) Click **Add** and then click **Apply**.
6. Repeat step 5 to create the default route IP prefix set with the following configuration.

Setting	Value
Name	Default Route
Network	0.0.0.0/0
Action	Permit

7. On the **Set IP prefix list** dialog box, click **Close**.

Configure Route Maps in the Tier-0 Gateway for Availability Zone 2

To define which routes are redistributed in the domain, you configure route maps in the tier-0 gateway.

1. On the NSX Manager main navigation bar, click **Networking**.
2. In the navigation pane, click **Tier-0 gateways**.
3. Select the gateway, and from the ellipsis menu, click **Edit**.
4. Create a route map for traffic incoming to availability zone 2.
 - a) Expand the **Routing** section and in the **Route maps** section, click **Set**.
 - b) In the **Set route maps** dialog box, click **Add route map**.
 - c) Enter a name for the route map.
For example, `rm-in-az2`.
 - d) In the **Match criteria** column, click **Set**.
 - e) On the **Set match criteria** dialog box, click **Add match criteria** and configure the following settings.

Setting	Value for Default Route	Value for Any
Type	IP Prefix	IP Prefix
Members	Default Route	Any
Local Preference	80	90
Action	Permit	Permit

- f) Click **Add** and then click **Apply**.
 - g) In the **Set route maps** dialog box, click **Save**.
5. Repeat step 4 to create a route map for outgoing traffic from availability zone 2 with the following configuration.

Setting	Value
Route map name	<code>rm-out-az2</code>
Type	IP Prefix
Members	Any
As Path Prepend	<code>bgp_asn</code>
Local Preference	100
Action	Permit

6. In the **Set route maps** dialog box, click **Close**.

Configure BGP in the Tier-0 Gateway for Availability Zone 2

To enable fail-over from availability zone 1 to availability zone 2, you configure BGP neighbors on the tier-0 gateway in the management or workload domain to be stretched. You add route filters to configure `localpref` on incoming traffic and `prepend of AS` on outgoing traffic.

You configure two BGP neighbors with route filters for the uplink interfaces in availability zone 2.

Table 233: BGP Neighbors for Availability Zone 2

Setting	BGP Neighbor 1	BGP Neighbor 2
IP address	<i>ip_bgp_neighbor1</i>	<i>ip_bgp_neighbor2</i>
BFD	Deactivated	Deactivated
Remote AS	<i>asn_bgp_neighbor1</i>	<i>asn_bgp_neighbor2</i>
Hold downtime	12	12
Keep alive time	4	4
Password	<i>bgp_password</i>	<i>bgp_password</i>

Table 234: Route Filters for BGP Neighbors for Availability Zone 2

Setting	BGP Neighbor 1	BGP Neighbor 2
IP Address Family	IPV4	IPV4
Activated	Activated	Activated
Out Filter	rm-out-az2	rm-out-az2
In Filter	rm-in-az2	rm-in-az2
Maximum Routes	-	-

1. On the NSX Manager main navigation bar, click **Networking**.
2. In the navigation pane, click **Tier-0 gateways**.
3. Select the gateway and from the ellipsis menu, click **Edit**.
4. Add the uplink interfaces to the NSX Edge nodes.
 - a) Expand **BGP** and in the **BGP neighbors** section, click **2**.
 - b) In the **Set BGP neighbors** dialog box, click **Add BGP neighbor** and configure the following settings.

Setting	Value
IP address	<i>ip_bgp_neighbor1</i>
BFD	Deactivated NOTE Activate BFD only if the network supports and is configured for BFD.
Remote AS	<i>asn_bgp_neighbor1</i>
Source addresses	Select AZ2 interfaces
Hold downtime	12
Keep alive time	4
Password	<i>bgp_password</i>

- c) In the **Route filter** section, click **Set**.
- d) In the Set route filter dialog box, click **Add route filter** and configure the following settings.

Setting	Value
IP Address Family	IPV4
Enabled	Activated
Out Filter	rm-out-az2
In Filter	rm-in-az2
Maximum Routes	-

- e) Click **Add** and then click **Apply**.
- Repeat step 4 to configure BGP neighbor `ip_bgp_neighbor2` and the corresponding route filter.
 - On the **Tier-0 gateway** page, click **Close editing**.

Configure Witness Traffic Separation for VMware Cloud Foundation on Dell VxRail

Witness traffic separation allows you to use a VMkernel adapter for vSAN witness traffic that is different from the adapter for vSAN data traffic.

You must have a stretched cluster before you can configure it for witness traffic separation.

By default, when you stretch a cluster, the vSAN-tagged VMkernel adapter is used to carry traffic destined for the vSAN witness host. With witness traffic separation, you can use a separately tagged VMkernel adapter instead of extending the vSAN data network to the witness host. This feature allows for a more flexible network configuration by allowing for separate networks for node-to-node and node-to-witness communication.

Create Distributed Port Groups for Witness Traffic

Create a distributed port group for each availability zone on the vSphere Distributed Switch.

- Log in to the vSphere Client.
- Click **Menu** › **Networking**.
- Right-click the vSphere distributed switch for the cluster and select **Distributed Port Group** › **New Distributed Port Group**.
- Enter a name for the port group for the first availability zone and click **Next**.
For example, `AZ1_WTS_PG`.
- Change the VLAN type to **VLAN** and enter a VLAN ID.
- Select **Customize default policies** and click **Next**.
- On the **Security** page, click **Next**.
- On the **Traffic shaping** page, click **Next**.
- On the **Teaming and failover** page, modify the failover order of the uplinks to match the existing failover order of the management traffic and click **Next**.
- On the **Monitoring** page, click **Next**.
- On the **Miscellaneous** page, click **Next**.
- On the **Ready to Complete** page, review your selections and click **Finish**.
- Repeat these steps for the second availability zone.

Delete Routes to the Witness Host

When you stretch a cluster, a route to the witness host is added to each ESXi host in the stretched cluster. You must delete these routes to use witness traffic separation.

- In a web browser, log in to the first ESXi host in the stretched cluster using the VMware Host Client.

2. In the navigation pane, click **Manage** and click the **Services** tab.
3. Select the **TSM-SSH** service and click **Start** if not started.
4. Open an SSH connection to the first ESXi host in the stretched cluster.
5. Log in as `root`.
6. Run the following command:

```
esxcli network ip route ipv4 list
```

The output returns something like:

Network	Netmask	Gateway	Interface	Source
default	0.0.0.0	172.18.15.1	vmk2	MANUAL
169.254.0.0	255.255.255.0	0.0.0.0	vmk1	MANUAL
172.18.7.0	255.255.255.0	0.0.0.0	vmk3	MANUAL
172.18.13.0	255.255.255.0	0.0.0.0	vmk5	MANUAL
172.18.14.0	255.255.255.0	172.18.7.253	vmk3	MANUAL
172.18.15.0	255.255.255.0	0.0.0.0	vmk2	MANUAL
172.18.21.0	255.255.255.0	172.18.7.253	vmk3	MANUAL

7. Delete the route to the witness host. For example:


```
esxconfig-route -d 172.18.14.0/24 172.18.7.253
```
8. In the VMware Host Client, select the **TSM-SSH** service for the ESXi host and click **Stop**.
9. Repeat these steps for each ESXi host in the stretched cluster.

Add VMkernel Adapters for Witness Traffic

Add VMkernel adapters for witness traffic to each availability zone's distributed port group.

1. Log in to the vSphere Client.
2. Click **Menu** > **Networking**.
3. Right-click the witness distributed port group for the first availability zone, for example, `AZ1_WTS_PG`, and select **Add VMkernel Adapters**.
4. Click **+ Attached Hosts**, select the availability zone 1 hosts from the list, and click **OK**.
5. Click **Next**.
6. Accept the default VMkernel port settings and click **Next**.

NOTE

Do not select any services.

7. Select **Use static IPv4 settings** and enter the IP addresses and the subnet mask to use for the witness traffic separation network.
8. Click **Next**.
9. Review your selections and click **Finish**.
10. Repeat these steps for the witness distributed port group for the second availability zone.

Configure the VMkernel Adapters for Witness Traffic

Enable witness traffic for the witness traffic VMkernel adapter on each ESXi host

1. Log in to the vSphere Client.
2. Click **Menu** > **Hosts and Clusters**.
3. For each host in the stretched cluster, click **Configure** > **Networking** > **VMkernel adapters** to determine which VMkernel adapter to use for witness traffic. For example, `vmk5`.
4. In a web browser, log in to the first ESXi host in the stretched cluster using the VMware Host Client.
5. In the navigation pane, click **Manage** and click the **Services** tab.
6. Select the **TSM-SSH** service and click **Start** if not started.
7. SSH to the first ESXi host in the stretched cluster.
8. Log in as root and run the following command:

```
esxcli vsan network ip add -i <vmkernel_adapter> -T=witness
```

For example:

```
esxcli vsan network ip add -i vmk5 -T=witness
```

9. Verify that the VMkernel adapter is configured for witness traffic:

```
esxcli vsan network list
```

10. Verify that the ESXi host can access the witness host:

```
vmkping -I <vmkernel_adapter><witness_host_ip_address>
```

Replace `<vmkernel_adapter>` with the VMkernel adapter configured for witness traffic, for example `vmk5`.

Replace `<witness_host_ip_address>` with the witness host IP address.

11. In the VMware Host Client, select the **TSM-SSH** service for the ESXi host and click **Stop**.
12. Repeat for each ESXi host in the stretched cluster.

Expand a Stretched VxRail Cluster

You can expand a stretched cluster by adding more VxRail nodes to the preferred and non-preferred sites.

You must have a stretched cluster.

1. Use the VxRail vCenter plug-in to add the additional hosts in availability zone 1 or availability zone 2 to the cluster by performing the VxRail Manager cluster expansion work flow.
Refer to the Dell VxRail documentation for more details.
2. Log in to SDDC Manager and run the script to trigger the workflow to import the newly added hosts in the SDDC Manager inventory.
In the script, provide the root credentials for each host and specify which fault domain the host should be added to.
3. Using SSH, log in to the SDDC Manager VM with the username `vcf` and the password you specified in the deployment parameter workbook.
4. Run the following command to expand the stretched cluster:

```
python initiate_stretch_cluster_vxrail.py --workflow expand-stretch-cluster --sc-domain <SDDC-valid-domain-name> --sc-cluster <valid cluster name which is a part of the domain to be stretched> --sc-hosts <valid host names> --witness-host-fqdn <witness host/appliance IP or fqdn> --witness-vsan-ip <witness vsan IP address> --witness-vsan-cidr <witness-vsan-network-IP-address-with-mask>
```

Replace `<SDDC-valid-domain-name>`, `<valid cluster name which is a part of the domain to be stretched>`, `<valid host names>`, `<witness vsan IP address>`, `<witness host/appliance IP or fqdn>`, `<witness vsan IP address>`, and `<witness-vsan-network-IP-address-with-mask>` with the correct values for your environment.

5. When prompted, enter the following information:
 - SSO user name and password
 - Root user password for ESXi hosts
 - Fault domain for ESXi hosts
 - vSAN gateway IP for the preferred (primary) and non-preferred (secondary) site
 - vSAN CIDR for the preferred (primary) and non-preferred (secondary) site
 - Confirm the SSH thumbprints for the hosts
6. Once the workflow is triggered, track the task status in the SDDC Manager UI.
If the task fails, debug and fix the issue and retry from SDDC Manager UI. Do not run the script again.

If you add hosts to a stretched cluster configured for witness traffic separation, perform the following tasks for the added hosts:

- [Add VMkernel Adapters for Witness Traffic](#)
- [Delete Routes to the Witness Host](#)
- [Configure the VMkernel Adapters for Witness Traffic](#)

Replace a Failed Host in a Stretched VxRail Cluster

If a host or host component in a stretched cluster fails, it is recommended that you replace the host with a new host.

- Check the health of the cluster.
See "Check vSAN Health" in *Administering VMware vSAN*.
1. Remove the failed host from the cluster.
See [Remove a Host from a Cluster in a Workload Domain](#).
 2. Expand the cluster to add the new host to the cluster.
See [Expand a Stretched VxRail Cluster](#).

vSAN automatically rebuilds the stretch cluster.

Monitoring Capabilities in the VMware Cloud Foundation System

The VMware Cloud Foundation system provides built-in capabilities to help you perform effective operations monitoring, troubleshooting, performance management, infrastructure capacity planning, and compliance monitoring and auditing.

You use the built-in monitoring capabilities for these typical scenarios.

Scenario	Examples
Are the systems online?	A host or other component shows a failed or unhealthy status.

Table continued on next page

Continued from previous page

Scenario	Examples
Why did a storage drive fail?	Hardware-centric views spanning inventory, configuration, usage, and event history to provide for diagnosis and resolution.
Is the infrastructure meeting tenant service level agreements (SLAs)?	Analysis of system and device-level metrics to identify causes and resolutions.
At what future time will the systems get overloaded?	Trend analysis of detailed system and device-level metrics, with summarized periodic reporting.
What person performed which action and when?	History of secured user actions, with periodic reporting. Workflow task history of actions performed in the system.

Viewing Tasks and Task Details

From SDDC Manager UI, you can access all tasks. By default, the Dashboard displays the Recent Tasks widget, providing general information at a glance about the most recent tasks. A task is a unit of work or a series of subtasks that perform an overall goal, such as creating a workload domain.

In addition to the most recent tasks, you can view and search for all tasks by clicking **View All Tasks** at the bottom of the Recent Tasks widget. This opens the Tasks panel.

NOTE

For more information about controlling the widgets that appear on the Dashboard page of SDDC Manager UI, see [Tour of the User Interface](#).

Viewing and Filtering Task Details

The Tasks panel provides a high level view all tasks, displaying the descriptive task name, task status (for example, running, succeeded, or failed), and the timestamp for the last change in task status. You can also filter and search the task information as follows:

- Search tasks by clicking the filter icon in the Task column header and entering a search string.
- Filter tasks by status by clicking the filter icon in Status column. Select by category **All**, **Failed**, **Successful**, **Running**, or **Pending**.

NOTE

Each category also displays the number of tasks with that status.

- Clear all filters by clicking **Reset Filter** at the top of the Tasks panel.
- Click **Refresh** to refresh the task list.

NOTE

You can also sort the table by the contents of the Status and Last Occurrence columns.

Managing Tasks and Subtask Details

Expand a task to view details including the subtasks that comprise the task and their individual statuses.

- If a task is in a Failed state, you can also attempt to restart it by clicking **Restart Task**.

NOTE

Not all tasks are restartable.

- If a task is in a Failed state, click on the icon next to the Failed status to view a detailed report on the cause.
- To view subtasks and their details, click **View Subtasks**.

NOTE

You can filter subtasks in the same way you filter tasks.

NOTE

You can also sort the table by the contents of the Status and Last Occurrence columns.

Resizing the Task Panel

Use the icons on the task panel to increase or decrease the panel size, or to close or reopen it.

API Activity Logging

When you invoke APIs or log in to or log out from the SDDC Manager UI, VMware Cloud Foundation creates activity log files that track the request. Activity logs can be used to analyze the pattern of user actions and gather metrics.

The following logs are available on the SDDC Manager appliance:

Log Name	Location
sddc-manager-ui-activity.log	/var/log/vmware/vcf/sddc-manager-ui-app
domainmanager-activity.log	/var/log/vmware/vcf/domainmanager
operationsmanager-activity.log	/var/log/vmware/vcf/operationsmanager
lcm-activity.log	/var/log/vmware/vcf/lcm
vcf-commonsvcs-activity.log	/var/log/vmware/vcf/commonsvcs

Activity Log Structure

All activity logs use the following JSON schema:

```
{
  "timestamp": "", "username": "", "clientIP": "", "userAgent": "", "api": "", "httpMethod": "",
  "httpStatus": "", "operation": "", "remoteIP": ""
}
```

Activity Log Example

The following example is from the domainmanager-activity.log:

```
{"username": "administrator@vsphere.local", "timestamp": "2022-01-19T16:59:01.9192 ",
"client IP": "10.0.0.253", "userAgent": "Mozilla/5.0 (Windows NT 6.3; Win 64; x64)
AppleWebKit/537.36 (KHTML, like Gecko) Chrome/97.0.4692.71 Safari/537.36", "api": "/
domainmanager/vl/vra/domains", "httpMethod": "GET", "httpStatus": 200, "operation": "Gets
VMware Aria Automationintegration status for workload domains", "remote IP": "127.0.0.1"}
```

- **username:** The username of the system from which the API request is triggered. For example: "administrator@vsphere.local".
- **timestamp:** Date and time of the operation performed in the UTC format "YYYY-MM-DD'TH:MM:SS.SSSXXX". For example: "2022-01-19T16:59:01.9192".
- **client IP:** The IP address of the user's system. For example: "10.0.0.253".
- **userAgent:** The user's system information such as the web browser name, web browser version, operating system name, and operating system architecture type. For example: "Mozilla/5.0 (Windows NT 6.3; Win 64; x64) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/97.0.4692.71 Safari/537.36".
- **api:** The API invoked to perform the operation. For example: "/domainmanager/vl/vra/domains".
- **httpMethod:** HTTP method of the REST API. For example: "GET".
- **httpStatus:** The response code received after invoking the API. For example: 200.

- **operation:** The operation or activity that was performed. For example: "Gets VMware Aria Automation integration status for workload domains".
- **remotelP:** remotelP of the request initiator. For example: "127.0.0.1"

Activity Logs Retention Policy

Log files are rolled over daily to a file using the following naming format: <service-name>.<YYYY>-<MM>-<DD>.0.log.gz. For example: domainmanager.2022-01-22.0.log.gz.

The log history is stored for 30 days. The maximum file size of the log retention file is set to 100 MB.

Log Analysis

You can perform log aggregation and analysis by integrating VMware Aria Operations for Logs with VMware Cloud Foundation. For more information, see [Implementation of Intelligent Logging and Analytics for VMware Cloud Foundation](#).

Updating VMware Cloud Foundation DNS and NTP Servers

If you need to update the DNS or NTP servers that VMware Cloud Foundation uses, you can update the servers using the SDDC Manager UI.

When you initially deploy VMware Cloud Foundation, you complete the deployment parameter workbook to provide the system with the information required for bring-up. This includes up to two DNS servers and up to two NTP servers. You can reconfigure these settings at a later date, using the SDDC Manager UI.

Update DNS Server Configuration

Use this procedure to update the DNS server configuration across VMware Cloud Foundation components.

- Verify that both forward and reverse DNS resolution are functional for each VMware Cloud Foundation component using the updated DNS server information.
- Verify that the new DNS server is reachable from each of the VMware Cloud Foundation components.
- Verify all VMware Cloud Foundation components are reachable from SDDC Manager.
- Verify that all VMware Cloud Foundation components are in an *Active* state.

SDDC Manager uses DNS servers to provide name resolution for the components in the system. When you update the DNS server configuration, SDDC Manager performs DNS configuration updates for the following components:

- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX Managers
- NSX Edge nodes
- VMware Aria Suite Lifecycle
- VMware Aria Operations for Logs
- VxRail Manager

If the update fails, SDDC Manager rolls back the DNS settings for the failed component. Fix the underlying issue and retry the update starting with the failed component.

NOTE

There is no rollback for VMware Aria Suite Lifecycle. Check the logs, resolve any issues, and retry the update.

Updating the DNS server configuration can take some time to complete, depending on the size of your environment. Schedule DNS updates at a time that minimizes the impact to the system users.

This procedure uses the SDDC Manager UI.

1. In the SDDC Manager UI, click **Administration > Network Settings**.
2. On the **Network Settings** page, click the **DNS Configuration** tab.
3. To update the DNS servers, click **Edit**.

4. Update the DNS configuration.
 - a) Expand the **Overview** section, and click **Next**.
 - b) Expand the **Prerequisites** section, and click **Next**.
 - c) Expand the **Edit DNS configuration** section, update the **Primary DNS server** and **Alternative DNS server**, and click **Save**.

NOTE

Alternative DNS server is optional.

Update NTP Server Configuration

Use this procedure to update the NTP server configuration across VMware Cloud Foundation components.

- Verify the new NTP server is reachable from the VMware Cloud Foundation components.
- Verify the time skew between the new NTP servers and the VMware Cloud Foundation components is less than 5 minutes.
- Verify all VMware Cloud Foundation components are reachable from SDDC Manager.
- Verify all VMware Cloud Foundation components are in an *Active* state.

SDDC Manager uses NTP servers to synchronize time between the components in the system. You must have at least one NTP server. When you update the NTP server configuration, SDDC Manager performs NTP configuration updates for the following components:

- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX Managers
- NSX Edge nodes
- VMware Aria Suite Lifecycle
- VMware Aria Operations for Logs
- VMware Aria Operations
- VMware Aria Automation
- VxRail Manager

If the update fails, SDDC Manager rolls back the NTP settings for the failed component. Fix the underlying issue and retry the update starting with the failed component.

NOTE

There is no rollback for the VMware Aria Suite Lifecycle. Check the logs, resolve any issues, and retry the update.

Updating the NTP server configuration can take some time to complete, depending on the size of your environment. Schedule NTP updates at a time that minimizes the impact to the system users. This procedure uses the SDDC Manager UI.

1. In the SDDC Manager UI, click **Administration > Network Settings**.
2. On the **Network Settings** page, click the **NTP Configuration** tab.
3. To update the NTP servers, click **Edit**.
4. Update the NTP configuration.
 - a) Expand the **Overview** section, and click **Next**.
 - b) Expand the **Prerequisites** section, and click **Next**.
 - c) Expand the **Edit NTP configuration** section, update the **NTP server**, and click **Save**.

Supportability and Serviceability (SoS) Utility

The SoS utility is a command-line tool that you can use to run health checks, collect logs for VMware Cloud Foundation components, and so on.

To run the SoS utility, SSH in to the SDDC Manager appliance using the `vcf` user account. For basic operations, enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-1 --option-2 --option-3 ... --option-n
```

To list the available command options, use the `--help` long option or the `-h` short option.

```
sudo /opt/vmware/sddc-support/sos --help
```

```
sudo /opt/vmware/sddc-support/sos -h
```

NOTE

You can specify options in the conventional GNU/POSIX syntax, using `--` for the long option and `-` for the short option.

For privileged operations, enter `su` to switch to the root user, and navigate to the `/opt/vmware/sddc-support` directory and type `./sos` followed by the options required for your desired operation.

SoS Utility Options

This section lists the specific options you can use with the SoS utility.

For information about collecting log files using the SoS utility, see [Collect Logs for Your System](#).

SoS Utility Help Options

Use these options to see information about the SoS utility itself. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

Option	Description
<code>--help</code> <code>-h</code>	Provides a summary of the available SoS utility options
<code>--version</code> <code>-v</code>	Provides the SoS utility's version number.

SoS Utility Generic Options

These are generic options for the SoS utility. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

Option	Description
<code>--history</code>	Displays the last 20 SoS operations performed.

Table continued on next page

Continued from previous page

Option	Description
<code>--force</code>	Allows SoS operations to be performed while workflows are running. NOTE It is recommended that you do not use this option.
<code>--configure-sftp</code>	Configures SFTP for logs.
<code>--setup-json SETUPJSON</code>	Custom setup-json file for log collection. SoS prepares the inventory automatically based on the environment where it is running. If you want to collect logs for a pre-defined set of components, you can create a <code>setup.json</code> file and pass the file as input to SoS. A sample JSON file is available on the SDDC Manager appliance at <code>/opt/vmware/sddc-support/setup.sample.json</code> .
<code>--log-folder LOGFOLDER</code>	Specifies the name of the log directory.
<code>--log-dir LOGDIR</code>	Specifies the directory to store the logs.
<code>--enable-stats</code>	Activate SoS execution stats collection.
<code>--debug-mode</code>	Runs the SoS utility in debug mode.
<code>--zip</code>	Creates a zipped TAR file for the output.
<code>--short</code>	Display detailed health results only for failures and warnings.
<code>--domain-name DOMAINNAME</code>	Specify the name of the workload domain name on which to perform the SoS operation. To run the operation on all workload domains, specify <code>--domain-name ALL</code> . NOTE If you omit the <code>--domain-name</code> flag and workload domain name, the SoS operation is performed only on the management domain. You can combine <code>--domain-name</code> with <code>--clusternames</code> to further limit the scope of an operation. This can be useful in a scaled environment with a large number of ESXi hosts.
<code>--clusternames CLUSTER NAMES</code>	Specify the vSphere cluster names associated with a workload domain for which you want to collect ESXi and Workload Management (WCP) logs. Enter a comma-separated list of vSphere clusters. For example, <code>--clusternames cluster1, cluster2</code> . NOTE If you specify <code>--domain-name ALL</code> then the <code>--clusternames</code> option is ignored.
<code>--skip-known-host-check</code>	Skips the specified check for SSL thumbprint for host in the known host.
<code>--include-free-hosts</code>	Collect logs for free ESXi hosts, in addition to in-use ESXi hosts.
<code>--include-precheck-report</code>	This option runs LCM upgrade prechecks and includes the LCM upgrade prechecks run report in SoS health check operations.

SoS Utility VMware Cloud Foundation Summary Options

These options provide summary details of the SDDC Manager instance, including components, services, and tasks.. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:


```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

Option	Description
<code>--get-vcf-summary</code>	Returns information about your VMware Cloud Foundation system, including CEIP, workload domains, vSphere clusters, ESXi hosts, licensing, network pools, SDDC Manager, and VCF services.
<code>--get-vcf-tasks-summary</code>	Returns information about VMware Cloud Foundation tasks, including the time the task was created and the status of the task.
<code>--get-vcf-services-summary</code>	Returns information about SDDC Manager uptime and when VMware Cloud Foundation services (for example, LCM) started and stopped.

SoS Utility Fix-It-Up Options

Use these options to manage ESXi hosts and vCenter Servers, including enabling SSH and locking down hosts. For these options, SSH in to the SDDC Manager VM using the `vcf` administrative user account, enter `su` to switch to the root user, navigate to the `/opt/vmware/sddc-support` directory, and type the following command:

```
./sos --option-name
```

NOTE

For Fix-It-Up options, if you do not specify a workload domain, the command affects only the management domain.

Option	Description
<code>--enable-ssh-esxi</code>	Applies SSH on all ESXi nodes in the specified workload domains. <ul style="list-style-type: none"> To enable SSH on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To enable SSH on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--disable-ssh-esxi</code>	Deactivates SSH on all ESXi nodes in the specified workload domains. <ul style="list-style-type: none"> To deactivate SSH on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To deactivate SSH on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--enable-ssh-vc</code>	Applies SSH on vCenter Server in the specified workload domains. <ul style="list-style-type: none"> To enable SSH on vCenter in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To enable SSH on vCenter Servers in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--disable-ssh-vc</code>	Deactivates SSH on vCenter Servers in the specified workload domains. <ul style="list-style-type: none"> To deactivate SSH on vCenter Server in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To deactivate SSH on vCenter Servers in all workload domains, include the flag <code>--domain-name ALL</code>.

Table continued on next page

Continued from previous page

Option	Description
<code>--enable-lockdown-esxi</code>	Applies normal lockdown mode on all ESXi nodes in the specified workload domains. <ul style="list-style-type: none"> To enable lockdown on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To enable lockdown on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--disable-lockdown-esxi</code>	Deactivates normal lockdown mode on ESXi nodes in the specified workload domains. <ul style="list-style-type: none"> To deactivate lockdown on ESXi nodes in a specific workload domain, include the flag <code>--domain-name DOMAINNAME</code>. To deactivate lockdown on ESXi nodes in all workload domains, include the flag <code>--domain-name ALL</code>.
<code>--ondemand-service ONDEMANDSERVICE</code>	Execute commands on ESXi hosts, vCenter Servers, or SDDC Manager entities for a given workload domain. Specify the workload domain using <code>--domain-name DOMAINNAME</code> . Replace <code>ONDEMANDSERVICE</code> with the path to a <code>.yaml</code> input file. (Sample file available at: <code>/opt/vmware/sddc-support/ondemand_command_sample.yaml</code>). WARNING Contact Broadcom Support before using this option.
<code>--ondemand-service JSON file path</code>	Include this flag to execute commands in the JSON format on all ESXi hosts in a workload domain. For example, <code>/opt/vmware/sddc-support/<JSON file name></code>
<code>--refresh-ssh-keys</code>	Refreshes the SSH keys.

SoS Utility Health Check Options

These SoS commands are used for checking the health status of various components or services, including connectivity, compute, storage, database, workload domains, and networks. For these options, SSH in to the SDDC Manager VM using the `vcf` user account and enter the following command:

```
sudo /opt/vmware/sddc-support/sos --option-name
```

Enter the `vcf` password when prompted.

A green status indicates that the health is normal, yellow provides a warning that attention might be required, and red (critical) indicates that the component needs immediate attention.

Option	Description
<code>--health-check</code>	Performs all available health checks. Can be combined with <code>--run-vsan-checks</code> . For example: <code>sudo /opt/vmware/sddc-support/sos --health-check --run-vsan-checks</code>
<code>--connectivity-health</code>	Performs connectivity checks and validations for SDDC resources (NSX Managers, ESXi hosts, vCenter Servers, and so on). This check performs a ping status check, SSH connectivity status check, and API connectivity check for SDDC resources.

Table continued on next page

Continued from previous page

Option	Description
<code>--services-health</code>	Performs a services health check to confirm whether services within the SDDC Manager (like Lifecycle Management Server) and vCenter Server are running.
<code>--compute-health</code>	Performs a compute health check, including ESXi host licenses, disk storage, disk partitions, and health status.
<code>--storage-health</code>	Performs a check on the vSAN disk health of the ESXi hosts and vSphere clusters. Can be combined with <code>--run-vsan-checks</code> . For example: <pre>sudo /opt/vmware/sddc-support/sos --storage-health --run-vsan-checks</pre>
<code>--run-vsan-checks</code>	This option cannot be run on its own and must be combined with <code>--health-check</code> or <code>--storage-health</code> . Runs a VM creation test to verify the vSAN cluster health. Running the test creates a virtual machine on each host in the vSAN cluster. The test creates a VM and deletes it. If the VM creation and deletion tasks are successful, assume that the vSAN cluster components are working as expected and the cluster is functional. NOTE You must not conduct the proactive test in a production environment as it creates network traffic and impacts the vSAN workload.
<code>--ntp-health</code>	Verifies whether the time on the components is synchronized with the NTP server in the SDDC Manager appliance. It also ensures that the hardware and software time stamp of ESXi hosts are within 5 minutes of the SDDC Manager appliance.
<code>--dns-health</code>	Performs a forward and reverse DNS health check.
<code>--general-health</code>	Checks ESXi for error dumps and gets NSX Manager and cluster status.
<code>--certificate-health</code>	Verifies that the component certificates are valid and when they are expiring. <ul style="list-style-type: none"> • GREEN: Certificate expires in more than 30 days. • YELLOW: Certificate expires in 15-30 days. • RED: Certificate expires in less than 15 days.
<code>--get-host-ips</code>	Returns host names and IP addresses of ESXi hosts.
<code>--get-inventory-info</code>	Returns inventory details for the VMware Cloud Foundation components, such as vCenter ServerNSX, SDDC Manager, and ESXi hosts. Optionally, add the flag <code>--domain-name ALL</code> to return details for all workload domains.
<code>--password-health</code>	Checks the status of passwords across VMware Cloud Foundation components. It lists components with passwords managed by VCF, the date a password was last changed, the password expiration date, and the number of days until expiration. <ul style="list-style-type: none"> • GREEN: Password expires in more than 15 days. • YELLOW: Password expires in 5-15 days. • RED: Password expires in less than 5 days.
<code>--hardware-compatibility-report</code>	Validates ESXi hosts and vSAN devices and exports the compatibility report.
<code>--version-health</code>	This operation checks the version of BOM components (vCenter Server, NSX, ESXi, and SDDC Manager). It compares the SDDC Manager inventory, the actual installed BOM component version, and the BOM component versions to detect any drift.

Table continued on next page

Continued from previous page

Option	Description
<code>--json-output-dir</code> JSONDIR	Outputs the results of any health check as a JSON file to the specified directory, JSONDIR.

Example Health Check Commands:

- Check the password health on the management domain only:
`./sos --password-health`
- Check the connectivity health for all workload domains:
`./sos --connectivity-health --domain-name ALL`
- Check the DNS health for the workload domain named `sfo-w01`:
`./sos --dns-health --domain-name sfo-w01`

Collect Logs for Your VMware Cloud Foundation System

Use the SoS utility to collect the logs for various software components in the system.

Use these options when retrieving support logs from your environment's various components.

- If you run the SoS utility from SDDC Manager without specifying any component-specific options, the SoS tool collects SDDC Manager, API, and VMware Cloud Foundation summary logs. To collect all logs, use the `--collect-all-logs` options.

NOTE

SoS log collection may time out after 60 minutes, which could be an issue with large workload domains. If the SoS utility does time out, collect component-specific logs or limit log collection to specific clusters using the options described below.

- If you run the SoS utility from Cloud Builder without specifying any component-specific options, the SoS tool collects SDDC Manager, API, and Cloud Builder logs.
- To collect logs for a specific component, run the utility with the appropriate options. For example, the `--domain-name` option is important. If omitted, the SoS operation is performed only on the management domain. See [SoS Utility Options](#).

After running the SoS utility, you can examine the resulting logs to troubleshoot issues, or provide to VMware Technical Support if requested. VMware Technical Support might request these logs to help resolve technical issues when you have submitted a support request. The diagnostic information collected using the SoS utility includes logs for the various VMware software components and software products deployed in your VMware Cloud Foundation environment.

Table 235: SoS Utility Log File Options

Option	Description
<code>--esx-logs</code>	Collects logs from the ESXi hosts only. Logs are collected from each ESXi host available in the deployment.
<code>--vc-logs</code>	Collects logs from the vCenter Server instances only. Logs are collected from each vCenter server available in the deployment.
<code>--sddc-manager-logs</code>	Collects logs from the SDDC Manager only. <code>sddc<timestamp>.tgz</code> contains logs from the SDDC Manager file system's <code>etc</code> , <code>tmp</code> , <code>usr</code> , and <code>var</code> partitions.

Table continued on next page

Continued from previous page

Option	Description
<code>--vxrail-manager-logs</code>	Collects logs from VxRail Manager instances only.
<code>--psc-logs</code>	Collects logs from the Platform Services Controller instances only.
<code>--nsx-logs</code>	Collects logs from the NSX Manager and NSX Edge instances only.
<code>--wcp-logs</code>	Collects logs from Workload Management clusters only.
<code>--vrealize-logs</code>	Collects logs from VMware Aria Suite Lifecycle.
<code>--no-clean-old-logs</code>	Use this option to prevent the utility from removing any output from a previous collection run. By default, before writing the output to the directory, the utility deletes the prior run's output files that might be present. If you want to retain the older output files, specify this option.
<code>--test</code>	Collects test logs by verifying the files.
<code>--no-health-check</code>	Skips the health check executed as part of log collection.
<code>--api-logs</code>	Collects output from REST endpoints for SDDC Manager inventory and LCM.
<code>--rvc-logs</code>	Collects logs from the Ruby vSphere Console (RVC) only. RVC is an interface for ESXi and vCenter. NOTE If the Bash shell is not enabled in vCenter Server, RVC log collection will be skipped . NOTE RVC logs are not collected by default with <code>./sos</code> log collection. You must enable RVC to collect RVC logs.
<code>--vm-screenshots</code>	Collects all VM screenshots.
<code>--system-debug-logs</code>	Collects system logs to help with debugging uncommon issues.
<code>--collect-all-logs</code>	Collects logs for all components, except Workload Management and system debug logs. By default, logs are collected for the management domain components. To collect logs for all workload domain, specify <code>--domain-name ALL</code> . To collect logs for a specific workload domain, specify <code>--domain-name domain_name</code> .
<code>--log-dir LOGDIR</code>	Specifies the directory to store the logs.
<code>--log-folder LOGFOLDER</code>	Specifies the name of the log directory.
<code>--domain-name DOMAINNAME</code>	Specify the name of the workload domain name on which the SoS operation is to be performed. To run the operation on all domains, specify <code>--domain-name ALL</code> . NOTE If you omit the <code>--domain-name</code> flag and domain name, the SoS operation is performed only on the management domain.

1. Using SSH, log in to the SDDC Manager appliance as the `vcf` user.
2. To collect the logs, run the SoS utility without specifying any component-specific options.

```
sudo /opt/vmware/sddc-support/sos
```

Enter the `vcf` password when prompted.

To collect logs for a specific component, run the utility with the appropriate options.

```
sudo /opt/vmware/sddc-support/sos --option-name
```

NOTE

By default, before writing the output to the directory, the utility deletes the prior run's output files that might be present. If you want to retain the older output files, specify the `--no-clean-old-logs` option.

If you do not specify the `--log-dir` option, the utility writes the output to the `/var/log/vmware/vcf/sddc-support` directory in the SDDC Manager appliance

The utility collects the log files from the various software components in all of the racks and writes the output to the directory named in the `--log-dir` option. Inside that directory, the utility generates output in a specific directory structure.

```
vcf@sddc-manager [ ~ ]$ sudo /opt/vmware/sddc-support/sos --domain-name MGMT --skip-known-host-check --log-dir /tmp/new
```

```
[sudo] password for vcf
```

```
Welcome to Supportability and Serviceability(SoS) utility!
```

```
Performing SoS operation for MGMT domain components
```

```
Logs : /tmp/new/sos-2019-09-03-21-04-40-11793
```

```
Log file : /tmp/new/sos-2019-09-03-21-04-40-11793/sos.log
```

```
Log Collection completed successfully for : [HEALTH-CHECK, SDDC-MANAGER, NSX_MANAGER, API-LOGS, ESX, VMS_SCREENSHOT, VCENTER-SERVER, VCF-SUMMARY]
```

Change to the output directory to examine the collected log files.

Component Log Files Collected by the SoS Utility

The SoS utility writes the component log files into an output directory structure within the file system of the SDDC Manager instance in which the command is initiated, for example:

```
vcf@sddc-manager [ ~ ]$ sudo /opt/vmware/sddc-support/sos
```

```
[sudo] password for vcf
```

```
Welcome to Supportability and Serviceability(SoS) utility!
```

```
Performing SoS operation for MGMT domain components
```

```
Logs : /var/log/vmware/vcf/sddc-support/sos-2019-09-03-20-55-41-10053
```

```
Log file : /var/log/vmware/vcf/sddc-support/sos-2019-09-03-20-55-41-10053/sos.log
```

```
NOTE : The Health check operation was invoked without --skip-known-host-check, and so will skip Connectivity Health, Password Health and Certificate Health Checks because of security reasons.
```

Log Collection completed successfully for : [HEALTH-CHECK, SDDC-MANAGER, NSX_MANAGER, API-LOGS, ESX, VMS_SCREENSHOT, VCENTER-SERVER, VCF-SUMMARY]

esx Directory Contents

In each rack-specific directory, the `esx` directory contains the following diagnostic files collected for each ESXi host in the rack:

File	Description
<code>esx-FQDN.tgz</code>	Diagnostic information from running the <code>vm-support</code> command on the ESXi host. An example file is <code>esx-esxi-1.vrack.vsphere.local.tgz</code> .
<code>SmartInfo-FQDN.txt</code>	S.M.A.R.T. status of the ESXi host's hard drive (Self-Monitoring, Analysis, and Reporting Technology). An example file is <code>SmartInfo-esxi-1.vrack.vsphere.local.txt</code> .
<code>vsan-health-FQDN.txt</code>	VMware vSAN cluster health information from running the standard command <code>python /usr/lib/vmware/vsan/bin/vsan-health-status.py</code> on the ESXi host. An example file is <code>vsan-health-esxi-1.vrack.vsphere.local.txt</code> .

nsx Directory Contents

In each rack-specific directory, the `nsx` directory contains the diagnostic information files collected for the NSX Managers and NSX Edge instances deployed in that rack.

The number of files in this directory depends on the number of NSX Manager and NSX Edge instances that are deployed in the rack. In a given rack, each management domain has a cluster of three NSX Managers. The first VI workload domain has an additional cluster of three NSX Managers. Subsequent VI workload domains can deploy their own NSX Manager cluster, or use the same cluster as an existing VI workload domain. NSX Edge instances are optional.

File	Description
<code>VMware-NSX-Manager-tech-support-<i>nsxmanagerIPaddr</i>.tar.gz</code>	Standard NSX Manager compressed support bundle, generated using the NSX API POST <code>https://<i>nsxmanagerIPaddr</i>/api/1.0/appliance-management/techsupportlogs/NSX</code> , where <code><i>nsxmanagerIPaddr</i></code> is the IP address of the NSX Manager instance. An example is <code>VMware-NSX-Manager-tech-support-10.0.0.8.tar.gz</code> .
<code>VMware-NSX-Edge-tech-support-<i>nsxmanagerIPaddr</i>-<i>edgeId</i>.tgz</code>	Standard NSX Edge support bundle, generated using the NSX API to query the NSX Edge support logs: GET <code>https://<i>nsxmanagerIPaddr</i>/api/4.0/edges/<i>edgeId</i>/techsupportlogs</code> , where <code><i>nsxmanagerIPaddr</i></code> is the IP address of the NSX

Table continued on next page

Continued from previous page

File	Description
<p>NOTE This information is only collected if NSX Edges are deployed.</p>	<p>Manager instance and <i>edgeID</i> identifies the NSX Edge instance. An example is VMware-NSX-Edge-tech-support-10.0.0.7-edge-1.log.gz.</p>

vc Directory Contents

In each rack-specific directory, the `vc` directory contains the diagnostic information files collected for the vCenter Server instances deployed in that rack.

The number of files in this directory depends on the number of vCenter Server instances that are deployed in the rack. In a given rack, each management domain has one vCenter Server instance, and any VI workload domains in the rack each have one vCenter Server instance.

File	Description
<code>vc-vcSaFQDN-vm-support.tgz</code>	Standard vCenter Server support bundle downloaded from the vCenter Server Appliance instance having a fully qualified domain name <i>vcSaFQDN</i> . The support bundle is obtained from the instance using the standard <code>vc-support.sh</code> command.

Managing Users and Groups in VMware Cloud Foundation

You can add users and groups to VMware Cloud Foundation to provide users with access to the SDDC Manager UI as well as the vCenter Server and NSX Manager instances that are deployed in your VMware Cloud Foundation system. Users can log in and perform tasks based on their assigned role.

Before you can add users and groups to VMware Cloud Foundation, you must configure an identity provider that has access to user and group data. VMware Cloud Foundation supports the following identity providers:

- vCenter Single Sign-On is vCenter Server's built-in identity provider. By default, it uses the system domain (for example, `vsphere.local`) as its identity source. You can add Active Directory over LDAP and OpenLDAP as identity sources for vCenter Single Sign-On.
- You can also use any of the following external identity providers instead of vCenter Single Sign-On:
 - Microsoft ADFS
 - Okta
 - Microsoft Entra ID (formerly known as Azure Active Directory)

Once you have configured an identity provider, you can add users and groups, and assign roles to determine what tasks they can perform from the SDDC Manager UI and VMware Cloud Foundation API.

NOTE

SDDC Manager only manages users and groups for the management SSO domain. If you created isolated VI workload domains that use different SSO domains, you must use the vSphere Client to manage users and groups for those SSO domains. Use the vSphere Client to connect to the VI workload domain's vCenter Server and then click **Administration > Single Sign On**.

In addition to user accounts, VMware Cloud Foundation includes the following accounts:

- Automation accounts for accessing VMware Cloud Foundation APIs. You can use these accounts in automation scripts.
- Local account for accessing VMware Cloud Foundation APIs when vCenter Server is down.

- Service accounts are automatically created by VMware Cloud Foundation for inter-product interaction. These are for system use only.

Configuring the Identity Provider for VMware Cloud Foundation

You can use vCenter Single Sign-On, Microsoft ADFS, Okta, or Microsoft Entra ID as the identity provider for VMware Cloud Foundation.

By default, VMware Cloud Foundation uses vCenter Single Sign-On as its identity provider and the system domain (for example, `vsphere.local`) as its identity source. You can add Active Directory over LDAP and OpenLDAP as identity sources for vCenter Single Sign-On. See [Add Active Directory over LDAP or OpenLDAP as an Identity Source for VMware Cloud Foundation](#).

You can also configure VMware Cloud Foundation to use Microsoft ADFS, Okta, or Microsoft Entra ID as an external identity provider, instead of using vCenter Single Sign-On:

- [Configure Microsoft ADFS as the Identity Provider in the SDDC Manager UI](#)
- [Configure Okta as the Identity Provider in the SDDC Manager UI](#)
- [Configure Identity Federation in VMware Cloud Foundation Using Microsoft Entra ID](#)

Add Active Directory over LDAP or OpenLDAP as an Identity Source for VMware Cloud Foundation

Users can log in to the SDDC Manager UI only if they are in a domain that has been added as a vCenter Single Sign-On identity source. vCenter Single Sign-On administrator users can add identity sources, or change the settings for identity sources that they added.

You can use identity sources to attach one or more domains to vCenter Single Sign-On. A domain is a repository for users and groups that the vCenter Single Sign-On server can use for user authentication with VMware Cloud Foundation. By default, vCenter Single Sign-On includes the system domain (for example, `vsphere.local`) as an identity source. You can add Active Directory over LDAP or an OpenLDAP directory service as identity sources.

1. In the navigation pane, click **Administration** > **Single Sign On**.
2. Click **Identity Provider**.
3. Click **Add** and select **AD over LDAP** or **OpenLDAP**.

Connect Identity Provider

> Overview Connecting your identity provider

2. AD over LDAP selected

Select the identity provider you want to assign to the vCenter server.

Select Identity Provider
Embedded

Embedded supports LDAP and Open LDAP. Add the identity source below.

Select Identity Source
AD over LDAP

NEXT

3. Server Settings Active Directory over LDAP settings

4. Review Review information and Submit

4. Click **Next**.
5. Enter the server settings and click **Next**.

Table 236: Active Directory over LDAP and OpenLDAP Server Settings

Option	Description
Identity Source Name	Name of the identity source.
Base Distinguished Name for Users	Base Distinguished Name for users. Enter the DN from which to start user searches. For example, <code>cn=Users,dc=myCorp,dc=com</code> .
Base Distinguished Name for Groups	The Base Distinguished Name for groups. Enter the DN from which to start group searches. For example, <code>cn=Groups,dc=myCorp,dc=com</code> .
Domain Name	The FQDN of the domain.
Domain Alias	For Active Directory identity sources, the domain's NetBIOS name. Add the NetBIOS name of the Active Directory domain as an alias of the identity source if you are using SSPI authentications. For OpenLDAP identity sources, the domain name in capital letters is added if you do not specify an alias.
User Name	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups. The ID can be in any of these formats: <ul style="list-style-type: none"> • UPN (user@domain.com)

Table continued on next page

Continued from previous page

Option	Description
	<ul style="list-style-type: none"> NetBIOS (DOMAIN\user) DN (cn=user,cn=Users,dc=domain,dc=com) The user name must be fully-qualified. An entry of "user" does not work.
Password	Password of the user who is specified by Username .
Primary Server URL	Primary domain controller LDAP server for the domain. You can use either the host name or the IP address. Use the format <code>ldap://hostname_or_IPaddress:port</code> or <code>ldaps://hostname_or_IPaddress:port</code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use <code>ldaps://</code> in the primary or the secondary LDAP URL.
Secondary Server URL	Address of a secondary domain controller LDAP server that is used for failover. You can use either the host name or the IP address.
Certificates (for LDAPS)	If you want to use LDAPS with your Active Directory LDAP Server or OpenLDAP Server identity source, click Browse to select a certificate. To export the root CA certificate from Active Directory, consult the Microsoft documentation.

6. Review the information and click **Submit**.

After you successfully add an identity source, you can add users and groups from the domain. See [Add a User or Group to](#) .

Configure Microsoft ADFS as the Identity Provider in the SDDC Manager UI

You can configure VMware Cloud Foundation to use Microsoft ADFS as an external identity provider, instead of using vCenter Single Sign-On. In this configuration, the external identity provider interacts with the identity source on behalf of vCenter Server.

Microsoft Active Directory Federation Services (ADFS) requirements:

- Microsoft ADFS for Windows Server 2016 or later must already be deployed.
- Microsoft ADFS must be connected to Active Directory.
- You have created a vCenter Server administrators group in Microsoft ADFS that contains the users you want to grant vCenter Server administrator privileges to.

For more information about configuring Microsoft ADFS, see the Microsoft documentation.

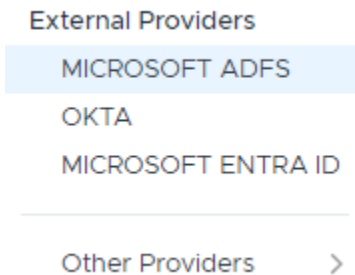
vCenter Server and other requirements:

- vSphere 7.0 or later

- vCenter Server must be able to connect to the Microsoft ADFS discovery endpoint, and the authorization, token, logout, JWKS, and any other endpoints advertised in the discovery endpoint metadata.

You can only add one external identity provider to VMware Cloud Foundation.

1. Log in to the SDDC Manager UI as a user with the ADMIN role
2. In the navigation pane, click **Administration** › **Single Sign On**.
3. Click **Identity Provider**.
4. Click **Change Identity Provider** and select **Microsoft ADFS**.



5. Click **Next**.
6. Select the checkbox to confirm the prerequisites and click **Next**.
7. If your Microsoft ADFS server certificate is signed by a publicly trusted Certificate Authority, click **Next**. If you are using a self-signed certificate, add the Microsoft ADFS root CA certificate added to the Trusted Root Certificates Store.
 - a) Click **Browse**.
 - b) Navigate to the certificate and click **Open**.
 - c) Click **Next**.
8. Copy the redirect URIs.

You will need them when you create the Microsoft ADFS Application Group in the next step.

9. Create an OpenID Connect configuration in Microsoft ADFS.

To establish a relying party trust between vCenter Server and an identity provider, you must establish the identifying information and a shared secret between them. In Microsoft ADFS, you do so by creating an OpenID Connect configuration known as an Application Group, which consists of a Server application and a Web API. The two components specify the information that vCenter Server uses to trust and communicate with the Microsoft ADFS server. To enable OpenID Connect in Microsoft ADFS, see the VMware knowledge base article at <https://kb.vmware.com/s/article/78029>.

Note the following when you create the Microsoft ADFS Application Group.

- You need the two Redirect URIs from the previous step.
 - Copy the following information to a file or write it down for use when configuring the identity provider in the next step.
 - Client Identifier
 - Shared Secret
 - OpenID address of the Microsoft ADFS server
10. Enter the Application Group information and click **Next**.
Use the information you gathered in the previous step and enter the:
 - Client Identifier
 - Shared Secret

- OpenID address of the Microsoft ADFS server

11. Enter user and group information for the Active Directory over LDAP connection to search for users and groups.

vCenter Server derives the AD domain to use for authorization and permissions from the Base Distinguished Name for users. You can add permissions on vSphere objects only for users and groups from this AD domain. Users or groups from AD child domains or other domains in the AD forest are not supported by vCenter Server Identity Provider Federation.

Option	Description
Base Distinguished Name for Users	Base Distinguished Name for users.
Base Distinguished Name for Groups	The base Distinguished Name for groups.
User Name	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Password	ID of a user in the domain who has a minimum of read-only access to Base DN for users and groups.
Primary Server URL	Primary domain controller LDAP server for the domain. Use the format <code>ldap://hostname:port</code> or <code>ldaps://hostname:port</code> . The port is typically 389 for LDAP connections and 636 for LDAPS connections. For Active Directory multi-domain controller deployments, the port is typically 3268 for LDAP and 3269 for LDAPS. A certificate that establishes trust for the LDAPS endpoint of the Active Directory server is required when you use <code>ldaps://</code> in the primary or secondary LDAP URL.
Secondary Server URL	Address of a secondary domain controller LDAP server that is used for failover.
Certificates (for LDAPS)	If you want to use LDAPS, click Browse to select a certificate.

12. Review the information and click **Submit**.

After you successfully add Microsoft ADFS as an external identity provider, you can add users and groups to VMware Cloud Foundation. See [Add a User or Group to](#) .


Add a User or Group to VMware Cloud Foundation

You can add users or groups so that they can access the SDDC Manager UI and VMware Cloud Foundation API.

Only a user with the ADMIN role can perform this task.

SDDC Manager UI displays user and group information based on the configured identity provider and identity sources. See [Configuring the Identity Provider for VMware Cloud Foundation](#).

1. In the navigation pane, click **Administration** > **Single Sign On**.
2. Click **Users and Groups** and then click **+ User or Group**.



3. Select one or more users or group by clicking the check box next to the user or group.

You can either search for a user or group by name, or filter by user type or domain.

4. Select a Role for each user and group.

Role	Description
ADMIN	This role has access to all the functionality of the UI and API.
OPERATOR	This role cannot access user management, password management, or backup configuration settings.
VIEWER	This role can only view the SDDC Manager. User management and password management are hidden from this role.

5. Scroll down to the bottom of the page and click **Add**.

Remove a User or Group

You can remove a user or group, for example when an employee leaves the company. The removed user or group will not be able to log in to the SDDC Manager UI.

Only a user with the ADMIN role can perform this task.

1. In the navigation pane, click **Administration > Single Sign On**.
2. Click the vertical ellipsis (three dots) next to a user or group name and click **Remove**.
3. Click **Delete**.

Create a Local Account

A local account is used to access VMware Cloud Foundation APIs when the management vCenter Server is down. If you upgraded from a previous release or didn't configure the account when deploying using the API, you can set a password using VMware Cloud Foundation APIs.

1. Log in to the SDDC Manager UI as a user with the ADMIN role.
For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).
2. In the navigation pane, click **Developer Center > API Explorer**.
3. To verify if the local account is configured, perform the following tasks:
 - a) Expand **APIs for managing Users**.
 - b) Expand `GET /v1/users/local/admin` and click **EXECUTE**.
 - c) In the Response, click `LocalUser (admin@local)`.

Response

```
LocalUser (admin@local) [🔗] ⬇ {  
  "isConfigured":  
    Flag indicating whether or not local  
    account is configured  
  "true",  
  "name":  
    The name of the user  
  "admin@local",  
  "role":  
    The role of the user  
  RoleReference (fa16f14b-9679-bbfc-06ed-47245405542c) [🔗] ⬇ { ... }  
  "type":  
    The type of the user  
  "USER",  
}
```

You can also download the response by clicking the download icon to the right of `LocalUser (admin@local)`.

4. If the local account is not configured, perform the following tasks to configure the local account:
 - a) Expand `PATCH /v1/users/local/admin`.
 - b) Enter a password for the local account and click **EXECUTE**.

PATCH /v1/users/local/admin Update password for local account

▼ Description
Update the password for local account only if the old password is correct, or if user configures the local account for the first time

> Response Types

▼ Try it out

Parameter	Value	Type	Description/Data Type
localUserPassword (required)	<pre>1 { 2 "newPassword": "Vmware123!Vmware123!" 3 }</pre>	Body	Local user password details LocalAccountPasswordInfo newPassword: (string, required) The new password of the local account oldPassword: (string) The old password of the local account }

EXECUTE COPY JSON DOWNLOAD

Password requirements are described below:

- Minimum length: 12
- Maximum length: 127
- At least one lowercase letter, one uppercase letter, a number, and one of the following special characters ! % @ \$ ^ # ? *
- A character cannot be repeated more than three times consecutively
- Must not include three of the same consecutive characters

NOTE

You must remember the password that you created because it cannot be retrieved. Local account passwords are used in password rotation.

Create an Automation Account

Automation accounts are used to access VMware Cloud Foundation APIs in automation scripts.

1. Log in to the SDDC Manager UI as a user with the ADMIN role.
For more about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).
2. In the navigation pane, click **Developer Center** > **API Explorer**.
3. Get the ID for the ADMIN role.
 - a) Expand **APIs for managing Users**.
 - b) Expand **GET /v1/roles** and click **Execute**.
 - c) In the Response, click **PageOfRole** and **Role (ADMIN)**.
 - d) Copy the ID for the ADMIN role.

Response

```

PageOfRole {
  "elements":
    The list of elements included in this page
    [
      Role (ADMIN) {
        "description":
          The description of the role
          "Administrator",
        "id":
          The ID of the role
          "317cb292-802f-ca6a-e57e-3ac2b707fe34",
        "name":
          The name of the role
          "ADMIN",
      },
    ]
  }
}

```

4. Create a service account with the ADMIN role and get the service account's API key.
 - a) Expand `POST /v1/users` and click **User**.
 - b) Replace the Value with:

```

[
  {
    "name": "service_account",
    "type": "SERVICE",
    "role":
      {
        "id": "317cb292-802f-ca6a-e57e-3ac2b707fe34"
      }
  }
]

```

Paste the ADMIN role ID from step 3.

POST /v1/users Add users

Description
Add list of users

Response Types

Try it out

Parameter	Value	Type	Description/Data Type
users (required)	<pre> 1 [2 { 3 "name": "service", 4 "role": { 5 "id": "317cb292-802f-ca6a-e57e-3ac2b70 6 }, 7 "type": "SERVICE" 8 } 9] </pre>	Body	User data collection [User{ ... }]

EXECUTE COPY JSON DOWNLOAD

- c) Click **Execute**.
- d) In the Response, click `PageOfUser` and `User (service_account)`.
- e) Copy the API key for the service account.

Response

`PageOfUser` {

"elements":
The list of elements included in this page
[

`User (service_account)` {

"apiKey":
The API key of the user
"qsfqnYgyxXQ892Jk9OHXyuEMgE3SgfTS",

5. Use the service account's API key to generate an access token.

- a) Expand **APIs for managing access and refresh tokens**.
- b) Expand **POST /v1/tokens**.
- c) Click **TokenCreationSpec**.
- d) Replace Value with:

```
{
  "apiKey": "qsfqnygyxXQ892Jk90HXyuEMgE3SgfTS"
}
```

Paste the service account's API key from step 4.

APIs for managing access and refresh token

POST /v1/tokens Create Token Pair

Description
Creates access token and refresh token for user access

Response Types

Try it out

Parameter	Value	Type	Description/Data Type
tokenCreationSpec (required)	<pre>{ 1: { 2: "apiKey": "qsfqnygyxXQ892Jk90HXyuEMgE3SgfTS" 3: } }</pre>	Body	tokenCreationSpec TokenCreationSpec{ ... }

EXECUTE COPY JSON DOWNLOAD

- e) Click **Execute**.
- f) In the Response, click **TokenPair** and **RefreshToken** and save the access and refresh tokens.

Response

TokenPair {

"accessToken":
Bearer token that can be used to make public API calls
"eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXZWQ892Jk90HXyuEMgE3SgfTS"

"refreshToken":
Refresh token that can be used to request new access token

RefreshToken (33f88c60-862e-4a38-8e8e-6479c4cd9f33) {

"id":
Refresh token id that can be used to request new access token
"33f88c60-862e-4a38-8e8e-6479c4cd9f33",

}

}

Managing Passwords in VMware Cloud Foundation

For security reasons, you can change passwords for the accounts that are used by your SDDC Manager instance. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

You entered passwords for your VMware Cloud Foundation system as part of the bring-up procedure. You can rotate and update some of these passwords using the password management functionality in the SDDC Manager UI, including:

- Accounts used for service consoles, such as the ESXi root account.
- The `root` and `mystic` users of the VxRail Manager
- The single sign-on administrator account(s).

NOTE

SDDC Manager manages passwords for all SSO administrator accounts, even if you created isolated VI workload domains that use different SSO domains than the management domain.

- The default administrative user account used by virtual appliances.
- Service accounts that are automatically generated during bring-up, host commissioning, and workload creation. Service accounts have a limited set of privileges and are created for communication between products. Passwords for service accounts are randomly generated by SDDC Manager. You cannot manually set a password for service accounts. To update the credentials of service accounts, you can rotate the passwords.

To provide optimal security and proactively prevent any passwords from expiring, you must rotate passwords every 80 days.

NOTE


Do not change the passwords for system accounts and the `administrator@vsphere.local` account outside SDDC Manager. This can break your VMware Cloud Foundation system.

You can also use the VMware Cloud Foundation API to look up and manage credentials. In the SDDC Manager UI, click **Developer Center** › **API Explorer** and browse to the APIs for managing credentials.

Starting with VMware Cloud Foundation 5.2.1, you can also manage passwords using the vSphere Client.


Password Expiration Notifications

The SDDC Manager UI provides a banner notification for any passwords managed by VMware Cloud Foundation that are expiring within the next 14 days. For example:



You can also click **Security** › **Password Management** in the navigation pane to view password expiration information. For example:

Expiring within 14 days ⓘ

 5 | **ESXI** 4 | **VCENTER** 1

[ROTATE ALL](#)

ESXI | **VCENTER** | **PSC** | **NSXT MANAGER** | **BACKUP**

[ROTATE ALL](#) | [ROTATE NOW](#) | [SCHEDULE ROTATION](#) ▾ | [RESET FILTERS](#)

<input type="checkbox"/>	User Name	FQDN	Status	Domain	Last Modified
<input type="checkbox"/>	root	esxi-1.vrack.vsphere.local	⚠ Expiring in 6 days	sddcid-1001	9/28/22, 1:33 AM
<input type="checkbox"/>	svc-vcf-esxi-1	esxi-1.vrack.vsphere.local	⚠ Expiring in 6 days	sddcid-1001	9/28/22, 1:33 AM
<input type="checkbox"/>	root	esxi-2.vrack.vsphere.local	⚠ Expiring in 6 days	sddcid-1001	9/28/22, 1:33 AM
<input type="checkbox"/>	svc-vcf-esxi-2	esxi-2.vrack.vsphere.local	⚠ Expiring in 6 days	sddcid-1001	9/28/22, 1:33 AM
<input type="checkbox"/>	svc-vcf-esxi-10	esxi-10.vrack.vsphere.local	✓ Active	UNASSIGNED	9/28/22, 2:15 AM

Expired passwords will display a status of **Disconnected**. For example:

ESXI | **vCenter** | **PSC** | **NSXT Manager** | **NSXT Edge** | **VRS LCM** | **Backup**

[ROTATE NOW](#) | [SCHEDULE ROTATION](#) ▾

<input type="checkbox"/>	User Name	FQDN	Domain	IP Address	Status
<input type="checkbox"/>	root	esx-7.vrack.vsphere.local	MGMT	10.0.0.101	🔴 Disconnected on 1/21/30, 12:00 AM.
<input type="checkbox"/>	root	esx-45.vrack.vsphere.local	MGMT	10.0.0.102	🔴 Disconnected on 1/19/30, 8:23 AM.
<input type="checkbox"/>	root	esx-qrs.vrack.vsphere.local	MGMT	10.0.0.103	🔴 Disconnected on 1/4/30, 8:36 AM.

For an expired password, you must update the password outside of VMware Cloud Foundation and then remediate the password using the SDDC Manager UI or the VMware Cloud Foundation API. See [Remediate Passwords](#).

NOTE

Password expiration information in the SDDC Manager UI is updated once a day. To get real-time information, use the VMware Cloud Foundation API.

Rotate Passwords

As a security measure, you can rotate passwords for the components in your VMware Cloud Foundation instance. The process of password rotation generates randomized passwords for the selected accounts. You can rotate passwords manually or set up auto-rotation for accounts managed by SDDC Manager.

- Verify that there are no currently failed workflows in SDDC Manager. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the brief time period that the password rotation process is running. It is recommended that you schedule password rotation for a time when you expect to have no running workflows.
- Only a user with the ADMIN role can perform this task.

You can rotate passwords for the following accounts.

- VxRail Manager
- ESXi

NOTE

Auto-rotate is not supported for ESXi.

- vCenter Server
By default, the vCenter Server root password expires after 90 days.

NOTE

Auto-rotate is automatically enabled for vCenter Server service accounts. It may take up to 24 hours to configure the service account auto-rotate policy for a newly deployed vCenter Server.

- vSphere Single-Sign On (PSC)
- NSX Edge nodes
- NSX Manager
- VMware Avi Load Balancer (formerly known as NSX Advanced Load Balancer)
- VMware Aria Suite Lifecycle
- VMware Aria Operations for Logs
- VMware Aria Operations
- VMware Aria Automation
- Workspace ONE Access

NOTE

For Workspace ONE Access passwords, the password rotation method varies depending on the user account. See the table below for details.

- SDDC Manager_{backup} user

Table 237: Password Rotation Details for Workspace ONE Access User Accounts

Workspace ONE Access User Account	VMware Aria Suite Lifecycle Locker Entry	Password Rotation Method	Password Rotation Scope
admin (443)	xint-wsa-admin	SDDC Manager Password Rotation	Application
admin (8443)	xint-wsa-admin	VMware Aria Suite Lifecycle Global Environment	Per node
configadmin (443)	xint-wsa-configadmin	<ol style="list-style-type: none"> 1. Reset the configadmin user password in Workspace ONE Access via the email reset link. 2. Create a new credential object in VMware Aria Suite Lifecycle Locker to match the new password. 3. Update the credential object referenced by <code>globalEnvironment</code> in VMware Aria Suite Lifecycle locker to the new credential object. 	Application

Table continued on next page

Continued from previous page

Workspace ONE Access User Account	VMware Aria Suite Lifecycle Locker Entry	Password Rotation Method	Password Rotation Scope
sshuser	global-env-admin	VMware Aria Suite Lifecycle Global Environment	Per node
root (ssh)	xint-wsa-root	SDDC Manager Password Rotation	Per node

The default password policy for rotated passwords requires:

- 20 characters in length
- At least one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ *
- No more than two of the same characters consecutively

If you changed the vCenter Server password length using the vSphere Client or the ESXi password length using the VMware Host Client, rotating the password for those components from SDDC Manager generates a password that complies with the password length that you specified.

To update the SDDC Manager root, super user, and API passwords, see [Updating SDDC Manager Passwords](#).

1. In the navigation pane, click **Security** › **Password Management**.
2. Select one or more accounts and click one of the following operation.
 - **Rotate Now**
 - **Schedule Rotation**
You can set the password rotation interval (30 days, 60 days, or 90 days). You can also deactivate the schedule.

NOTE

The **Schedule Rotation** option is not available for ESXi.

NOTE

Auto-rotate schedule is configured to run at midnight on the scheduled date. If auto-rotate could not start due to any technical issue, there is a provision to auto-retry every hour till start of the next day. In case schedule rotation is missed due to technical issues the UI displays a global notification with failed task status. The status of the schedule rotation can also be checked on the Tasks panel.

A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status for the password rotation operation. To view sub-tasks, click the task name. As each of these tasks is run, the status is updated. If the task fails, you can click **Retry**.

Password rotation is complete when all sub-tasks are completed successfully.

Manually Update Passwords

You can manually change the password for a selected account. Unlike password rotation, which generates a randomized password, you provide the new password.

- Verify that there are no currently failed workflows in your VMware Cloud Foundation system. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no active workflows are running or are scheduled to run during the manual password update.
- Only a user with the ADMIN role can perform this task. For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).

You can update only one password at a time.

Although individual VMware Cloud Foundation accounts support different password requirements, it is recommended that you set passwords following a common set of requirements across all accounts:

- Minimum length: 15
- Maximum length: 20
- At least one lowercase letter, one uppercase letter, a number, and one of the following special characters: ! @ # \$ ^ *
- Must NOT include:
 - A dictionary word
 - A palindrome
 - More than four monotonic character sequences
 - Three of the same consecutive characters

1. From the navigation pane, select **Security › Password Management**.
2. Select the account whose password you want to update, click the vertical ellipsis (three dots), and click **Update Password**.
3. Enter and confirm the new password.
4. Click **Update**.

A message appears at the top of the page showing the progress of the operation. The Tasks panel also shows detailed status of the password update operation. To view sub-tasks, click the task name.

Password update is complete when all sub-tasks are completed successfully.

Remediate Passwords

When an error occurs, for example after a password expires, you must manually reset the password in the component product. After you reset the password in a component, you must remediate the password in SDDC Manager to update the password in the SDDC Manager database and the dependent VMware Cloud Foundation workflows.

- Verify that VMware Cloud Foundation system contain no failed workflows. To check for failed workflows, click **Dashboard** in the navigation pane and expand the **Tasks** pane at the bottom of the page.
- Verify that no workflows are running or are scheduled to run while you remediate the password.
- Only a user with the ADMIN role can perform this task. For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).

To resolve any errors that might have occurred during password rotation or update, you must use password remediation. Password remediation syncs the password of the account stored in the SDDC Manager with the updated password in the component.

NOTE

You can remediate the password for only one account at a time.

Although the individual VMware Cloud Foundation components support different password requirements, you must set passwords following a common set of requirements across all components. For information on updating passwords manually, see [Manually Update Passwords](#).

1. From the navigation pane, select **Security › Password Management**.
2. Select the account whose password you want to remediate, click the vertical ellipsis (three dots), and click **Remediate Password**.

The Remediate Password dialog box appears. This dialog box displays the entity name, account type, credential type, and user name, in case you must confirm you have selected the correct account.

3. Enter and confirm the password that was set manually on the component.
4. Click **Remediate**.

A message appears at the top of the page showing the progress of the operation. The Task panel also shows detailed status of the password remediation operation. To view subtasks, you can click the task name.

Password remediation is complete when all sub-tasks are completed successfully.

Look Up Account Credentials

To look up the account credentials for the built-in accounts that are managed and rotated by SDDC Manager, you can log in to the SDDC Manager appliance using any SDDC Manager account credentials.

Only a user with the `ADMIN` role can perform this task.

1. SSH in to the SDDC Manager appliance using the `vcf` user account.
2. Change to the `/usr/bin` directory.

NOTE

Although the password management CLI commands are located in `/usr/bin`, you can run them from any directory.

3. Obtain the account credentials list by typing the command:

```
lookup_passwords
```

You must enter the user name and password for a user with the `ADMIN` role.

NOTE

Accounts with type `USER` and `SYSTEM` will be displayed.

4. Save the command output to a secure location with encryption so that you can access it later and use it to log in to the accounts as needed.

Updating SDDC Manager Passwords

The process for updating SDDC Manager passwords varies, depending on which account you are updating.

Update SDDC Manager Root and Super User Passwords

For security reasons, you can change passwords for the SDDC Manager root (`root`) and super user (`vcf`) accounts. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

The SDDC Manager `root` password expires after 90 days.

1. SSH in to the SDDC Manager VM using the `vcf` user account.
2. Enter `su` to switch to the root user.
3. Enter one of the following commands:

<code>passwd vcf</code>	To change the super user password.
<code>passwd root</code>	To change the root password.

4. Enter and retype the new password. For example:

```
root@sddc-manager [ /home/vcf ]# passwd vcf
```

```
New password:
```

```
Retype new password:
passwd: password updated successfully
```

The password is updated.

Update SDDC Manager Local Account Password

The SDDC Manager local account is used to access VMware Cloud Foundation APIs when the management vCenter Server is down. For security reasons, you should periodically update the password for this account.

Password requirements for the SDDC Manager local account:

- At least 15 characters
- No more than 127 characters
- At least one lowercase letter
- At least one uppercase letter
- At least one digit
- At least one special character, such as @ ! # \$ % ^ or ?
- A character cannot be repeated more than 3 times consecutively

1. Log in to the SDDC Manager UI as a user with the ADMIN role.

For more information about roles, see [Managing Users and Groups in VMware Cloud Foundation](#).

2. Click **Developer Center > API Explorer**.
3. Expand **APIs for managing Users**.
4. Expand `PATCH /v1/users/local/admin`.
5. In the **Description/Data Type** column, click `LocalAccountPasswordInfo{...}`.
6. In the **Value** box, type the new and old passwords and click **Execute**.
7. Click **Continue** to confirm.

A response of `Status: 204, No Content` indicates that the password was successfully updated.

Update Expired SDDC Manager Root Password

This section describes the procedure for updating an expired password for the SDDC Manager root (`root`) user.

The password must meet the following requirements:

- Minimum length 15 characters
- Must include:
 - mix of uppercase and lowercase letters
 - a number
 - a special character, such as @ ! # \$ % ^ or ?
- Must not include:
 - * { } [] () / \ ' " ` ~ , ; : . < >
 - A dictionary word (for example, `VMware1!`)

1. In a web browser, log in to the management domain vCenter Server using the vSphere Client (`https://<vcenter_server_fqdn>/ui`).
2. In the VMs and Templates inventory, expand the management domain vCenter Server and the management virtual machines folder.

3. Right-click the SDDC Manager virtual machine, and select **Open Remote Console**.
4. Click within the console window and press **Enter** on the Login menu item.
5. Type **root** as the user name and enter the current password for the root user.
6. Type `passwd root`.
7. When prompted for a new password, enter a different password than the previous one and click **Enter**.

Backing Up and Restoring SDDC Manager and NSX Manager

Regular backups of the management VMs are important to avoid downtime and data loss in case of a system failure. If a VM does fail, you can restore it to the last backup.

You can backup and restore SDDC Manager with an image-based or a file-based solution. File-based backup is recommended for customers who are comfortable with configuring backups using APIs, and are not using composable servers.

For a file-based backup of SDDC Manager VM, the state of the VM is exported to a file that is stored in a domain different than the one where the product is running. You can configure a backup schedule for the SDDC Manager VM and enable task-based (state-change driven) backups. When task-based backups are enabled, a backup is triggered after each SDDC Manager task (such as workload domain and host operations or password rotation).

You can also define a backup retention policy to comply with your company's retention policy. For more information, see the *VMware Cloud Foundation on Dell VxRail API Reference Guide*.

By default, NSX Manager file-based backups are taken on the SFTP server that is built into SDDC Manager. It is recommended that you configure an external SFTP server as a backup location for the following reasons:

- An external SFTP server is a prerequisite for restoring SDDC Manager file-based backups.
- Using an external SFTP server provides better protection against failures because it decouples NSX backups from SDDC Manager backups.

This section of the documentation provides instructions on backing up and restoring SDDC Manager, and on configuring the built-in automation of NSX backups. For information on backing up and restoring a full-stack SDDC, see *VMware Validated Design Backup and Restore*.

Reconfigure SFTP Backups for SDDC Manager and NSX Manager

By default, backups of SDDC Manager and NSX Manager are stored in the SDDC Manager appliance. Change the destination of the backups to an external SFTP server.

- Only a user with the ADMIN role can perform this task. See [Managing Users and Groups in VMware Cloud Foundation](#).
- The external SFTP server must support a 256-bit length ECDSA SSH public key.
- The external SFTP server must support a 2048-bit length RSA SSH public key
- You will need the SHA256 fingerprint of RSA key of the SFTP server.
- Host Key algorithms: At least one of `rsa-sha2-512` or `rsa-sha2-256` and one of `ecdsa-sha2-nistp256`, `ecdsa-sha2-nistp384`, or `ecdsa-sha2-nistp521`.
- Additional pre-requisites when FIPS Security Mode is enabled on SDDC Manager:

Algorithms and Ciphers	Required when FIPS Security Mode is Enabled
Kex Algorithms	At least one of: <ul style="list-style-type: none"> • <code>diffie-hellman-group-exchange-sha256</code> • <code>ecdh-sha2-nistp256</code> • <code>ecdh-sha2-nistp384</code> • <code>ecdh-sha2-nistp521</code>

Table continued on next page

Continued from previous page

Algorithms and Ciphers	Required when FIPS Security Mode is Enabled
Message Authentication Key (MAC) Algorithms	hmac-sha2-256
Ciphers	<p>At least one of:</p> <ul style="list-style-type: none"> • TLS_DHE_DSS_WITH_AES_128_CBC_SHA256 • TLS_DHE_DSS_WITH_AES_256_CBC_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_128_CBC_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384 • TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 • TLS_RSA_WITH_AES_128_CBC_SHA256 • TLS_RSA_WITH_AES_256_CBC_SHA256 • TLS_ECDHE_ECDSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 • TLS_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 • TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 • TLS_AES_128_GCM_SHA256 • TLS_AES_256_GCM_SHA384

NOTE

SHA1 algorithms are not supported.

1. In the navigation pane, click **Administration** > **Backup**.
2. On the **Backup** page, click the **Site Settings** tab and then click **Register External**.
3. On the **Backup** page, enter the settings and click **Save**.

To obtain the SSH Fingerprint of the target system to verify, connect to the SDDC Manager Appliance over ssh and run the following command:

```
ssh-keygen -lf <(ssh-keyscan -p 22 -t rsa sftp_server_fqdn 2> /dev/null) | cut -d' ' -f2
```

Setting	Value
Host FQDN or IP	The FQDN or IP Address of the SFTP server.
Port	22
Transfer Protocol	SFTP
Username	A service account with privileges to the SFTP server. For example: svc-vcf-bck.
Password	The password for the username provided.

Table continued on next page

Continued from previous page

Setting	Value
Backup Directory	The directory on the SFTP server where backups are saved. For example: /backups/.
SSH Fingerprint	The SSH Fingerprint is automatically retrieved from the SFTP server, verify the SSH Fingerprint.
Confirm Fingerprint	Selected
Encryption Passphrase	The encryption passphrase used to encrypt the backup data. NOTE The encryption passphrase should be stored safely as it is required during the restore process.

4. In the **Confirm your changes to backup settings** dialog box, click **Confirm**.

File-Based Backups for SDDC Manager and vCenter Server

You can use the native file-based backup capabilities of SDDC Manager, vCenter Server, and NSX Manager. The NSX Manager backup is configured by SDDC Manager during the bring-up process. You configure the file-based backup jobs for SDDC Manager and vCenter Server.

Verify that you have an SFTP server on the network to serve as a target of the file-based backups.

To ensure that all management components are backed up correctly, you must create a series of backup jobs that capture the state of a set of related components at a common point in time. With some components, simultaneous backups of the component nodes ensure that you can restore the component a state where the nodes are logically consistent with each other and eliminate the necessity for further logical integrity remediation of the component.

Table 238: File-Based Backup Jobs

Component	Recommended Frequency	Recommended Retention	Notes
SDDC Manager	Daily	7 days	You must configure the backup jobs for the SDDC Manager instance and all vCenter Server instances in the vCenter Single Sign-On domain to start within the same 5-minute window.
vCenter Server	Daily	7 days	
vSphere Distributed Switch	On-demand	Retain last 3 configurations.	-
NSX Manager	Hourly	7 days	Configured by SDDC Manager during the bring-up process.

NOTE

- You must monitor the space utilization on the SFTP server to ensure that you have sufficient storage space to accommodate all backups taken within the retention period.
- Do not make any changes to the `/opt/vmware/vcf` directory on the SDDC Manager VM. If this directory contains any large files, backups may fail.

Back Up SDDC Manager

You configure file-based daily backups of the SDDC Manager instances using the SDDC Manager administration interface.

Only a user with the **Admin** role can perform this task.

1. In the navigation pane, click **Administration > Backup**.
2. On the **Backup** page, click the **SDDC Manager Configurations** tab.
3. Under **Backup Schedule**, click **Edit**.
4. On the **Backup Schedule** page, enter the settings and click **Save**.

Setting	Value
Automatic Backup	Enabled
Backup Frequency	Weekly
Days of the Week	All selected
Schedule Time	04:02 AM
Take Backup on State Change	Enabled
Retain Last Backups	7
Retain Hourly Backups for Days	1
Retain Daily Backups for Days	7

5. To verify the backup, click **Backup Now**.

The status and the start time of the backup is displayed on the UI. You have set the SDDC Manager backup schedule to run daily at 04:02 AM and after each change of state.

If the backup is unsuccessful, verify if the SFTP server is available and able to provide its SSH fingerprint:

- SSH to the SDDC Manager appliance run the following command as the root user:

```
sftp username@IP of sftp server
```

Enter the SFTP user password when prompted. The following message indicates a successful connection:

```
Connected to username@IP of sftp server.
```

- To check that the SFTP server SSH fingerprint is available, run:

```
ssh-keygen -lf <(ssh-keyscan -t ssh-rsa -p port_numberIP of sftp server 2>/dev/null)
```

Configure a Backup Schedule for vCenter Server

You configure file-based daily backups of the vCenter Server instances by using the vCenter Server Management Interface of each vCenter Server instance.

1. In a web browser, log in to the vCenter Server Management Interface (`https://appliance-IP-address-or-FQDN:5480`).
2. In the left navigation pane, click **Backup**.

3. In the **Backup schedule** pane, click **Configure**.
4. In the **Create backup schedule** dialog box, enter these values and click **Create**.

Setting		Value
Backup location		Enter the backup location from SFTP server. For example: <code>sftp://172.16.11.60/backups/</code>
Backup server credentials	User name	A service account with privileges to the SFTP server. For example: <code>svc-vcf-bck</code> .
	Password	Enter the password for the username provided.
Schedule		Daily 11:00 PM
Encrypt backup	Encryption password	<code>encryption_password</code>
	Confirm password	<code>encryption_password</code>
Number of backups to retain		Retain last 7 backups
Data	Stats, events, and tasks	Selected
	Inventory and configuration	Selected

The backup schedule information appears in the **Backup schedule** pane.

5. Repeat the procedure for the other vCenter Server instances.

Any complete and in-progress backup appears in the **Activity** pane.

Manually Back Up vCenter Server

Before you upgrade a vCenter Server instance, you should use the vCenter Server Management Interface to manually back it up.

- In the vSphere Client, for each vSphere cluster that is managed by the vCenter Server, note the current vSphere DRS Automation Level setting and then change the setting to **Manual**. After the vCenter Server upgrade is complete, you can change the vSphere DRS Automation Level setting back to its original value. See [KB 87631](#) for information about using VMware PowerCLI to change the vSphere DRS Automation Level.
 - Ensure that there are not any active vMotion tasks.
1. In a web browser, log in to the vCenter Server Management Interface (`https://appliance-IP-address-or-FQDN:5480`).
 2. In the left navigation pane, click **Backup**.
 3. Click **Backup Now**.
 4. If you already have a backup schedule set up, select **Use backup location and user name from backup schedule** and click **Start**.
 5. If you do not already have a backup schedule, enter the following information and click **Start**.

Setting		Value
Backup location		Enter the backup location from SFTP server. For example: <code>sftp://172.16.11.60/backups/</code>
Backup server credentials	User name	A service account with privileges to the SFTP server. For example: <code>svc-vcf-bck</code> .
	Password	Enter the password for the username provided.
Encrypt backup	Encryption password	<code>encryption_password</code>
	Confirm password	<code>encryption_password</code>
Data	Stats, events, and tasks	Selected
	Inventory and configuration	Selected

In order to restore vCenter Server, you will need the VMware vCenter Server Appliance ISO file that matches the version you backed up.

- Identify the required vCenter Server version. In the vCenter Server Management Interface, click **Summary** in the left navigation pane to see the vCenter Server version and build number.
- Download the VMware vCenter Server Appliance ISO file for that version from the Broadcom Support Portal.

Export the Configuration of the vSphere Distributed Switches

The vCenter Server backup includes the configuration of the entire vCenter Server instance. To have a backup only of the vSphere Distributed Switch and distributed port group configurations, you export a configuration file that includes the validated network configurations. If you want to recover only the vSphere Distributed Switch, you can import this configuration file to the vCenter Server instance.

You can use the exported file to create multiple copies of the vSphere Distributed Switch configuration on an existing deployment, or overwrite the settings of existing vSphere Distributed Switch instances and port groups. You must backup the configuration of a vSphere Distributed Switch immediately after each change in configuration of that switch.

1. In a web browser, log in to vCenter Server by using the vSphere Client.
2. Select **Menu** > **Networking**.
3. In the inventory expand **vCenter Server** > **Datacenter**.
4. Expand the **Management Networks** folder, right-click the distributed switch, and select **Settings** > **Export configuration**.
5. In the **Export configuration** dialog box, select **Distributed switch and all port groups**.
6. In the **Description** text box enter the date and time of export, and click **OK**.
7. Copy the backup zip file to a secure location from where you can retrieve the file and use it if a failure of the appliance occurs.
8. Repeat the procedure for the other vSphere Distributed Switches.

File-Based Restore for SDDC Manager, vCenter Server, and NSX

When SDDC Manager, vCenter Server, or NSX Manager in the SDDC fails, you can restore the component to a fully operational state by using its file-based backup. When an NSX Edge node fails, you redeploy the node from the NSX Manager instance.

Use this guidance as appropriate based on the exact nature of the failure encountered within your environment. Sometimes, you can recover localized logical failures by restoring individual components. In more severe cases, such as a complete and irretrievable hardware failure, to restore the operational status of your SDDC, you must perform a complex set of manual deployments and restore sequences. In failure scenarios where there is a risk of data loss, there has already been data loss or where it involves a catastrophic failure, contact Broadcom Support to review your recovery plan before taking any steps to remediate the situation.

Restore SDDC Manager

If SDDC Manager fails, you can restore it from its file-based backup.

- Power off and rename the failed SDDC Manager instance.
- Verify that you have a valid file-based backup of the failed SDDC Manager instance.

To be valid, the backup must be of the same version as the version of the SDDC Manager appliance on which you plan to restore the instance.

- Verify that you have the SFTP server details:
 - SFTP Server IP
 - SFTP Server Username
 - SFTP Server Password
 - Encryption Password

After a successful recovery, securely delete the decrypted backup files.

Prepare for Restoring SDDC Manager

Before restoring SDDC Manager, you must download and decrypt the encrypted backup file from the SFTP server.

Verify that your host machine with access to the SDDC has OpenSSL installed.

NOTE

The procedures have been written based on the host machine being a Linux-based operating system.

The backup file contains sensitive data about your VMware Cloud Foundation instance, including passwords in plain text. As a best practice, you must control access to the decrypted files and securely delete them after you complete the restore operation.

1. Identify the backup file for the restore and download it from the SFTP server to your host machine.
2. On your host machine, open a terminal and run the following command to extract the content of the backup file.

```
OPENSSL_FIPS=1 openssl enc -d -aes-256-cbc -md sha256 -in filename-of-restore-file |
tar -xz
```

3. When prompted, enter the *encryption_password*.
4. In the extracted folder, locate and open the *metadata.json* file in a text editor.
5. Locate the *sddc_manager_ova_location* value and copy the URL.
6. In a web browser, paste the URL and download the OVA file.
7. In the extracted folder, locate and view the contents of the *security_password_vault.json* file.
8. Locate the *entityType BACKUP* value and record the backup password.

Restore SDDC Manager from a File-Based Backup

First, you deploy a new SDDC Manager appliance by using the OVA file that you downloaded during the preparation for the restore. After that, you restore the file-based backup on the newly deployed SDDC Manager appliance.

1. In a web browser, log in to management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > VMs and templates**.
3. In the inventory expand **vCenter Server > Datacenter**.
4. Right-click the management folder and select **Deploy OVF template**.
5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, browse to the location of the SDDC Manager OVA file, click **Open**, and click **Next**.
6. On the **Select a name and folder** page, in the **Virtual machine name** text box, enter a virtual machine name, and click **Next**.
7. On the **Select a compute resource** page, click **Next**.
8. On the **Review details** page, review the settings and click **Next**.
9. On the **License agreements** page, accept the license agreement and click **Next**.
10. On the **Select storage** page, select the vSAN datastore and click **Next**.

The datastore must match the `vsan_datastore` value in the `metadata.json` file that you downloaded during the preparation for the restore.

11. On the **Select networks** page, from the **Destination network** drop-down menu, select the management network distributed port group and click **Next**.

The distributed port group must match the `port_group` value in the `metadata.json` file that you downloaded during the preparation for the restore.

12. On the **Customize template** page, enter the following values and click **Next**.

Setting	Description
Enter root user password	You can use the original root user password or a new password.
Enter login (vcf) user password	You can use the original vcf user password or a new password.
Enter basic auth user password	You can use the original admin user password or a new password.
Enter backup (backup) user password	The backup password that you saved during the preparation for the restore. This password can be changed later if desired.
Enter Local user password	You can use the original Local user password or a new password.
Hostname	The FQDN must match the <code>hostname</code> value in the <code>metadata.json</code> file that you downloaded during the preparation for the restore.
NTP sources	The NTP server details for the appliance.
Enable FIPs	Selected
Default gateway	The default gateway for the appliance.
Domain name	The domain name for the appliance.
Domain search path	The domain search path(s) for the appliance.

Table continued on next page

Continued from previous page

Setting	Description
Domain name servers	The DNS servers for the appliance.
Network 1 IP address	The IP address for the appliance.
Network 1 netmask	The subnet mask for the appliance.

13. On the **Ready to complete** page, click **Finish** and wait for the process to complete.
14. When the SDDC Manager appliance deployment completes, expand the management folder.
15. Right-click the SDDC Manager appliance and select **Snapshots > Take Snapshot**.
16. Right-click the SDDC Manager appliance, select **Power > Power On**.
17. On the host machine, copy the encrypted backup file to the `/tmp` folder on the newly deployed SDDC Manager appliance by running the following command. When prompted, enter the `vcf_user_password`.

```
scp filename-of-restore-file vcf@sddc_manager_fqdn:/tmp/
```

18. On the host machine, obtain the authentication token from the SDDC Manager appliance in order to be able to execute the restore process by running the following command:

```
TOKEN=`curl https://<sddc_manager_fqdn>/v1/tokens -k -X POST -H "Content-Type: application/json" -d '{"username": "admin@local", "password": "<admin@local_password>" }' | awk -F "\"" '{ print $4}'`
```

19. On the host machine with access to the SDDC Manager, open a terminal and run the command to start the restore process.

```
curl https://<sddc_manager_fqdn>/v1/restores/tasks -k -X POST -H "Content-Type: application/json" -H "Authorization: Bearer $TOKEN" \
  -d '{
  "elements" : [ {
    "resourceType" : "SDDC_MANAGER"
  } ],
  "backupFile" : "<backup_file>",
  "encryption" : {
    "passphrase" : "<encryption_password>"
  }
}'
```

The command output contains the ID of the restore task.

20. Record the ID of the restore task.
21. Monitor the restore task by using the following command until the status becomes `Successful`.

```
curl https://<sddc_manager_fqdn>/v1/restores/tasks/<restore_task_id> -k -X GET -H "Content-Type: application/json" -H "Authorization: Bearer $TOKEN"
```

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Validate the Status of SDDC Manager

After a successful restore of SDDC Manager, you must validate its status. You run the health checks by using the `sos` tool.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client.
2. Run the health checks by using the `sos` tool.

```
sudo /opt/vmware/sddc-support/sos --health-check
```

3. When prompted, enter the `vcf_password`.
All tests show green when SDDC Manager is in healthy state.
4. Manually delete the snapshot created in [Restore SDDC Manager from a File-Based Backup](#).

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Restore vCenter Server

If a vCenter Server instance fails, you can restore it from its file-based backup.

- Power off and rename the failed vCenter Server instance.
- Verify that you have a valid file-based backup of the failed vCenter Server instance.

To be valid, the backup must be of the version of the vCenter Server Appliance on which you plan to restore the instance.

- Verify that you have the SFTP server details:
 - SFTP Server IP
 - SFTP Server Username
 - SFTP Server Password
 - Encryption Password

Prepare for Restoring vCenter Server

Before restoring a vCenter Server instance, you must retrieve the vCenter Server build number and deployment details, as well as vCenter Server and ESXi credentials from the SDDC Manager inventory.

SDDC Manager must be available.

Retrieve the vCenter Server Deployment Details

Before restoring a vCenter Server instance, you must retrieve the vCenter Server build number and deployment details from the SDDC Manager inventory. The vCenter Server instances in your system might be running different build numbers if the backups are taken during an upgrade process. You must restore each vCenter Server instance to its correct version.

Because the Management domain vCenter Server might be unavailable to authenticate the login, you use the SDDC Manager API via the shell to retrieve this information.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client.
2. Run the command to get the list of vCenter Server instances.

```
curl http://localhost/inventory/vcenters -k | json_pp
```

3. For each vCenter Server instance, record the values of these settings.

Setting	Value
domainType	Name of the domain
vmName	VM name of the vCenter Server
managementIpAddress	IP address of the vCenter Server
datastoreForVmDeploymentName	Datastore name
hostName	FQDN of the vCenter Server
version	<i>version_number-build_number</i>
Size	Size of the deployment

4. Verify that the vCenter Server version retrieved from SDDC Manager is the same as the version associated with the backup file that you plan to restore.

Retrieve the Credentials for Restoring vCenter Server

Before restoring a vCenter Server instance, you must retrieve the vCenter Server root and vCenter Single Sign-On administrator credentials from the SDDC Manager inventory. Before restoring the Management domain vCenter Server, you must also retrieve the credentials of a healthy Management domain ESXi host.

NOTE

If SDDC Manager is not operational, you can retrieve the required vCenter Server root, vCenter Single Sign-On administrator, and ESXi root credentials from the file-based backup of SDDC Manager. See [Prepare for Restoring SDDC Manager](#).

Before you can query the SDDC Manager API, you must obtain an API access token by using **admin@local** account.

1. Log in to your host machine with access to the SDDC and open a terminal.
2. Obtain the API access token.
 - a) Run the command to obtain an access token by using the **admin@local** credentials.

```
TOKEN=`curl https://<sddc_manager_fqdn>/v1/tokens -k -X POST -H "Content-Type: application/json" -d '{"username": "admin@local", "password": "admin@local_password"}' | awk -F "\"" '{print $4}'`
```

The command returns an access token and a refresh token.

- b) Record the access token.
3. Retrieve the vCenter Server **root** credentials.
 - a) Run the following command to retrieve the vCenter Server **root** credentials.

```
curl https://<sddc_manager_fqdn>/v1/credentials?resourceType=VCENTER -k -X GET \
-H "Accept: application/json" -H "Authorization: Bearer $TOKEN" | json_pp
```

The command returns the vCenter Server **root** credentials.

Setting	Value
domainName	Name of the domain
resourceName	FQDN of the vCenter Server
username	root

Table continued on next page

Continued from previous page

Setting	Value
password	<i>vcenter_server_root_password</i>

- b) Record the vCenter Server **root** credentials.
4. Retrieve the vCenter Single Sign-On administrator credentials.

- a) Run the following command to retrieve the vCenter Single Sign-On administrator credentials.

```
curl https://<sddc_manager_fqdn>/v1/credentials?resourceType=PSC -k -X GET \-H
"Accept: application/json" -H "Authorization: Bearer $TOKEN" | json_pp
```

The command returns the **administrator@vsphere.local** credentials.

Setting	Value
domainName	Name of the domain
resourceName	FQDN of the vCenter Server
username	administrator@vsphere.local
password	<i>vsphere_admin_password</i>

- b) Record the **administrator@vsphere.local** credentials.
5. If you plan to restore the management domain vCenter Server, retrieve the credentials for a healthy management domain ESXi host.
- a) Run the following command to retrieve the credentials for a management domain ESXi host.

```
curl https://<sddc_manager_fqdn>/v1/credentials?resourceType=ESXI -k -X GET \-H
"Accept: application/json" -H "Authorization: Bearer $TOKEN" | json_pp
```

The command returns the ESXi **root** credentials.

Setting	Value for first ESXi host
domainName	management domain name
resourceName	FQDN of the first ESXi host
username	root
password	<i>esxi_root_password</i>

- b) Record the ESXi **root** credentials.

Restore a vCenter Server Instance from a File-Based Backup

If a vCenter Server instance fails, you can restore it from its file-based backup. If the management domain vCenter Server and the VI workload domain vCenter Server are both in a failed state, you must restore the management domain vCenter Server before restoring the VI workload domain vCenter Server.

- Download the vCenter Server ISO file for the version of the failed instance. See [Retrieve the vCenter Server Deployment Details](#).
- If you are recovering the VI workload domain vCenter Server, verify that the management vCenter Server is available.

You deploy a new vCenter Server appliance and perform a file-based restore. If you are restoring the management domain vCenter Server, you deploy the new appliance on a healthy ESXi host in the management domain vSAN cluster. If you are restoring the VI workload domain vCenter Server, you deploy the new appliance on the management domain vCenter Server.

1. Mount the vCenter Server ISO image to your host machine with access to the SDDC and run the UI installer for your operating system.
For example, for a Windows host machine, open the `dvd-drive:\vcsa-ui-installer\win32\installer` application file.
2. Click **Restore**.
3. Complete the **Restore - Stage 1: Deploy vCenter Server** wizard.
 - a) On the **Introduction** page, click **Next**.
 - b) On the **End user license agreement** page, select the **I accept the terms of the license agreement** check box and click **Next**.
 - c) On the **Enter backup details** page, enter these values and click **Next**.

Setting	Value for vCenter Server
Location or IP/hostname	<code>sftp://sftp_server_ip/backups/vCenter/sn_vc_fqdn/backup_folder/</code>
User name	vSphere service account user
Password	<code>vsphere-service-account-password</code>

- d) On the **Review backup information** page, review the backup details, record the **vCenter Server configuration** information, and click **Next**.
You use the vCenter Server configuration information at a later step to determine the deployment size for the new vCenter Server appliance.
- e) On the **vCenter Server deployment target** page, enter the values by using the information that you retrieved during the preparation for the restore, and click **Next**.

Setting	Value for Management Domain vCenter Server	Value for VI Workload Domain vCenter Server
ESXi host or vCenter Server name	FQDN of the first ESXi host	FQDN of the management vCenter Server
HTTPS port	443	443
User name	root	administrator@vsphere.local
Password	<code>esxi_root_password</code>	<code>vsphere_admin_password</code>

- f) In the **Certificate warning** dialog box, click **Yes** to accept the host certificate.
- g) On the **Set up a target vCenter Server VM** page, enter the values by using the information that you retrieved during the preparation for the restore, and click **Next**.

Setting	Value
VM name	vCenter Server VM name
Set root password	<code>vcenter_server_root_password</code>
Confirm root password	<code>vcenter_server_root_password</code>

- h) On the **Select deployment size** page, select the deployment size that corresponds with the vCenter Server configuration information from Step 3.d and click **Next**.
Refer to vSphere documentation to map CPU count recorded from Step 3.d to a vSphere Server configuration size.
- i) On the **Select datastore** page, select these values, and click **Next**.

Setting	Value
Datastore	Datastore name
Enable thin disk mode	Selected

- j) On the **Configure network settings** page, enter the values by using the information that you retrieved during the preparation for the restore, and click **Next**.

Setting	Value
Network	Name of the vSphere distributed switch
IP version	IPv4
IP assignment	static
FQDN	FQDN of the vCenter Server
IP address	IP address of the vCenter Server
Subnet mask or prefix length	24
Default gateway	Default gateway IP address
DNS servers	DNS server IP addresses with comma separated

- k) On the **Ready to complete stage 1** page, review the restore settings and click **Finish**.
 l) When stage 1 of the restore process completes, click **Continue**.
4. Complete the **Restore - Stage 2: vCenter Server** wizard.
- a) On the **Introduction** page, click **Next**.
 b) On the **Backup details** page, in the **Encryption password** text box, enter the encryption password of the SFTP server and click **Next**.
 c) On the **Single Sign-On configuration** page, enter these values and click **Next**.

Setting	Value
Single Sign-On user name	administrator@vsphere.local
Single Sign-On password	<i>vsphere_admin_password</i>

- d) On the **Ready to complete** page, review the restore details and click **Finish**.
 e) In the **Warning** dialog box, click **OK** to confirm the restore.
 f) When stage 2 of the restore process completes, click **Close**.

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Move the Restored vCenter Server Appliance to the Correct Folder

After deploying and restoring a vCenter Server instance, you must move the new appliance to the correct folder.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > VMs and Templates**.
3. In the inventory expand **vCenter Server > Datacenter**.
4. Right-click the appliance of the restored vCenter Server instance and select **Move to folder**.

5. Select the management folder and click **OK**.

Validate the vCenter Server State

After restoring a vCenter Server instance, you must validate the state of the vCenter Server and vCenter Single Sign-On.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. In the inventory, click the management domain vCenter Server inventory, click the **Summary** tab, and verify that there are no unexpected vCenter Server alerts.
3. Click the **Linked vCenter Server systems** tab and verify that the list contains all other vCenter Server instances in the vCenter Single Sign-On domain.
4. Log in to the recovered vCenter Server instance by using a Secure Shell (SSH) client.
5. Run the command to navigate to the `bin` directory.

```
cd /usr/lib/vmware-vmware/bin
```

6. Validate the current replication status.
 - a) Run the command to list the current replication partners of the vCenter Server instance with the current replication status between the nodes.


```
./vdcadmin -f showpartnerstatus -h localhost -u administrator -w vsphere_admin_password
```
 - b) Verify that for each partner, the `vdcadmin` command output contains `Host available: Yes, Status available: Yes, and Partner is 0 changes behind`.
 - c) If you observe significant differences, because the resyncing might take some time, wait five minutes and repeat this step.
7. Repeat the procedure for the other vCenter Server instance.

Validate the SDDC Manager State After a vCenter Server Restore

After a successful vCenter Server restore, verify that the SDDC Manager inventory is consistent with the recovered VMs and that the vCenter Server instances are healthy. You use the Supportability and Serviceability tool (SoS) and the SDDC Manager patch/upgrade precheck function.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client.
2. Run the SoS health check and verify the output.

```
sudo /opt/vmware/sddc-support/sos --health-check
```

All tests show green when SDDC Manager is in a healthy state.

3. In a Web browser, log in to SDDC Manager using the user interface.
4. In the navigation pane, click **Inventory > Workload Domains**.
5. For each workload domain, validate the vCenter Server status.
 - a) Click the workload domain name and click the **Updates/Patches** tab.
 - b) Click **Precheck**.
 - c) Click **View status** to review the precheck result for the vCenter Server instance and verify that the status is `Succeeded`.

Restore the Configuration of a vSphere Distributed Switch

To recover the configuration of a vSphere Distributed Switch, you can restore its settings from the configuration file that you previously exported.

This procedure restores only the vSphere Distributed Switch configuration of a vCenter Server instance. The restore operation changes the settings on the vSphere Distributed Switch back to the settings saved in the configuration file. The operation overwrites the current settings of the vSphere Distributed Switch and its port groups. The operation does not delete existing port groups that are not a part of the configuration file. The vSphere Distributed Switch configuration is part of the vCenter Server backup. If you want to restore the entire vCenter Server instance, see [Restore vCenter Server](#).

1. In a web browser, log in to the vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** > **Networking**.
3. In the inventory expand **vCenter Server** > **Datacenter**.
4. Expand the **Management networks** folder, right-click the distributed switch and select **Settings** > **Restore configuration**.
5. On the **Restore switch configuration** page, click **Browse**, navigate to the location of the configuration file for the distributed switch, and click **Open**.
6. Select the **Restore distributed switch and all port groups** radio-button and click **Next**.
7. On the **Ready to complete** page, review the changes and click **Finish**.
8. Repeat these steps for the other vSphere Distributed Switch.
9. Review the switch configuration to verify that it is as you expect after the restore.

Restore an NSX Manager Cluster Node

If an NSX Manager instance fails, you can restore it from its file-based backup.

- Verify that you have a valid file-based backup of the failed NSX Manager instance.
- Verify that you have the SFTP server details:
 - SFTP Server IP
 - SFTP Server Username
 - SFTP Server Password
 - Encryption Password

Prepare for Restoring an NSX Manager Cluster Node

Before restoring an NSX Manager node, you must retrieve the NSX Manager build number and deployment details, as well as the credentials from the SDDC Manager inventory.

Retrieve the NSX Manager Version from SDDC Manager

Before restoring a failed NSX Manager instance, you must retrieve its version from the SDDC Manager inventory.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. Click the domain name of the failed NSX Manager instance.
3. Click the **Update/Patches** tab.
4. Under **Current versions**, in the **NSX** panel, locate and record the **NSX upgrade coordinator** value.
5. Verify that the NSX version retrieved from SDDC Manager is the same as the version associated with the backup file that you plan to restore.

Retrieve the Credentials for Restoring NSX Manager from SDDC Manager

Before restoring a failed NSX Manager instance, you must retrieve the NSX Manager **root** and **admin** credentials from the SDDC Manager inventory.

Before you can query the SDDC Manager API, you must obtain an API access token by using an API service account.

1. Log in to your host machine with access to the SDDC and open a terminal.

2. Obtain the API access token.

- a) Run the command to obtain an access token by using the **admin@local** account credentials.

```
curl 'https://<sddc_manager_fqdn>/v1/tokens' -k -X POST -H 'Content-Type: application/json' -H 'Accept: application/json' -d '{"username" : "service_user", "password" : "service_user_password"}'
```

The command returns an access token and a refresh token.

- b) Record the access token.

3. Retrieve the NSX Manager **root** and **admin** credentials.

- a) Run the command to retrieve the NSX Manager **root** and **admin** credentials.

```
curl 'https://<sddc_manager_fqdn>/v1/credentials?resourceType=NSXT_MANAGER' -i -X GET -H 'Accept: application/json' -H 'Authorization: Bearer access_token'
```

The command returns the NSX Manager **root** and **admin** credentials.

- b) Record the NSX Manager **root** and **admin** credentials for the instance you are restoring.

Restore the First Node of a Failed NSX Manager Cluster

If all three NSX Manager nodes in an NSX Manager cluster are in a failed state, you begin the restore process by restoring the first cluster node.

IMPORTANT

This procedure is not applicable in use cases when there are operational NSX Manager cluster nodes.

- If two of the three NSX Manager nodes in the NSX Manager cluster are in a failed state, you begin the restore process by deactivating the cluster. See [Deactivate the NSX Manager Cluster](#).
- If only one of the three NSX Manager nodes in the NSX Manager cluster is in a failed state, you directly restore the failed node to the cluster. See [Restore an NSX Manager Node to an Existing NSX Manager Cluster](#).

Redeploy the First Node of a Failed NSX Manager Cluster

You deploy a new NSX Manager instance by using the configuration of the first NSX Manager cluster node.

- Download the NSX Manager OVA file for the version of the failed NSX Manager cluster. See [Retrieve the NSX Manager Version from SDDC Manager](#).
 - Verify that the backup file that you plan to restore is associated with the version of the failed NSX Manager cluster.
1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
 2. Select **Menu** > **VMs and Templates**.
 3. In the inventory, expand **vCenter Server** > **Datacenter**.
 4. Right-click the NSX folder and select **Deploy OVF Template**.
 5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, navigate to the location of the NSX Manager OVA file, click **Open**, and click **Next**.
 6. On the **Select a name and folder** page, enter the VM name and click **Next**.
 7. On the **Select a compute resource** page, click **Next**.
 8. On the **Review details** page, click **Next**.
 9. On the **Configuration** page, select the appropriate size and click **Next**.

For the management domain, select **Medium** and for workload domains, select **Large** unless you changed these defaults during deployment.

10. On the **Select storage** page, select the vSAN datastore, and click **Next**.
11. On the **Select networks** page, from the **Destination network** drop-down menu, select the management network distributed port group, and click **Next**.
12. On the **Customize template** page, enter these values and click **Next**.

Setting	Value for first NSX Manager cluster node
System root user password	<i>nsx_root_password</i>
CLI admin user password	<i>nsx_admin_password</i>
CLI audit user password	<i>nsx_audit_password</i>
Hostname	Enter hostname for the appliance using FQDN format.
Default IPv4 gateway	Enter the default gateway for the appliance.
Management network IPv4 address	Enter the IP Address for the appliance.
Management network netmask	Enter the subnet mask for the appliance.
DNS server list	Enter the DNS servers for the appliance.
NTP server list	Enter the NTP server for the appliance.
Enable SSH	Deselected
Allow root SSH logins	Selected

13. On the **Ready to complete** page, review the deployment details and click **Finish**.

Restore the First Node in a Failed NSX Manager Cluster from a File-Based Backup

You restore the file-based backup of the first NSX Manager cluster node to the newly deployed NSX Manager instance.

1. In a web browser, log in to the NSX Manager node for the domain by using the user interface (https://<nsx_manager_node_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left navigation pane, under **Lifecycle management**, click **Backup and restore**.
4. In the **NSX configuration** pane, under **SFTP server**, click **Edit**.
5. In the **Backup configuration** dialog box, enter these values, and click **Save**.

Setting	Value
FQDN or IP address	IP address of SFTP server
Protocol	SFTP
Port	22
Directory path	/backups
Username	Service account user name For example, svc-vcf-bck@rainpole.io
Password	<i>service_account_password</i>
SSH fingerprint	<i>SFTP_ssh_fingerprint</i>

6. Under **Backup history**, select the target backup, and click **Restore**.
7. During the restore, when prompted, reject adding NSX Manager nodes by clicking **I understand** and **Resume**.

A progress bar displays the status of the restore operation with the current step of the process.

Validate the Status of the First NSX Manager Cluster Node

After you restored the first NSX Manager cluster node, you validate the services state from the VM Web console of the restored node.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **VMs and Templates**.
3. In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
4. Click the VM name of the newly deployed first NSX Manager cluster node, click **Launch Web Console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

5. Run the command to view the cluster status.

```
get cluster status
```

The services on the single-node NSX Manager cluster appear as UP.

Deactivate the NSX Manager Cluster

If two of the three NSX Manager cluster nodes are in a failed state or if you restored the first node of a failed NSX Manager cluster, you must deactivate the cluster.

IMPORTANT

This procedure is not applicable in use cases when there are two operational NSX Manager cluster nodes.

If only one of the three NSX Manager nodes in the NSX Manager cluster is in a failed state, after you prepared for the restore, you directly restore the failed node to the cluster. See [Restore an NSX Manager Node to an Existing NSX Manager Cluster](#).

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **VMs and Templates**.
3. In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
4. Click the VM of the operational NSX Manager node in the cluster, click **Launch Web Console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

5. Run the command to deactivate the cluster

```
deactivate cluster
```

- On the **Are you sure you want to remove all other nodes from this cluster? (yes/no)** prompt, enter `yes`.
You deactivated the cluster.

Power off and delete the two failed NSX Manager nodes from inventory.

Restore an NSX Manager Node to an Existing NSX Manager Cluster

If only one of the three NSX Manager cluster nodes is in a failed state, you restore the failed node to the existing cluster. If two of the three NSX Manager cluster nodes are in a failed state, you repeat this process for each of the failed nodes.

Detach the Failed NSX Manager Node from the NSX Manager Cluster

Before you recover a failed NSX Manager node, you must detach the failed node from the NSX Manager cluster.

- In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
- Select **Menu** › **VMs and Templates**.
- In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
- Click the VM of an operational NSX Manager node in the cluster, click **Launch Web Console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<code>nsx_admin_password</code>

- Retrieve the UUID of the failed NSX Manager node.
 - Run the command to view the details of the cluster members.

```
get cluster status
```

The status of the failed node is `Down`.

- Record the UUID of the failed NSX Manager node.
- Run the command to detach the failed node from the cluster

```
detach node failed_node_uuid
```

The detach process might take some time.

- When the detaching process finishes, run the command to view the cluster status.

```
get cluster status
```

The status of all cluster nodes is `Up`.

Redeploy the Failed NSX Manager Node

You deploy a new NSX Manager instance by using the configuration of the failed node.

Download the NSX Manager OVA file for the version of the failed NSX Manager instance. See [Retrieve the NSX Manager Version from SDDC Manager](#).

- In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
- Select **Menu** › **VMs and Templates**.

3. In the inventory expand **vCenter Server** › **Datacenter**.
4. Right-click the NSX folder and select **Deploy OVF Template**.
5. On the **Select an OVF template** page, select **Local file**, click **Upload files**, navigate to the location of the NSX Manager OVA file, click **Open**, and click **Next**.
6. On the **Select a name and folder** page, in the **Virtual machine name** text box, enter VM name of the failed node, and click **Next**.
7. On the **Select a compute resource** page, click **Next**.
8. On the **Review details** page, click **Next**.
9. On the **Configuration** page, select **Medium**, and click **Next**.
10. On the **Select storage** page, select the vSAN datastore, and click **Next**.
11. On the **Select networks** page, from the **Destination network** drop-down menu, select the management network distributed port group, and click **Next**.
12. On the **Customize template** page, enter these values and click **Next**.

Setting	Value
System root user password	<i>nsx_root_password</i>
CLI admin user password	<i>nsx_admin_password</i>
CLI audit password	<i>nsx_audit_password</i>
Hostname	<i>failed_node_FQDN</i>
Default IPv4 gateway	Enter the default gateway for the appliance.
Management network IPv4 address	<i>failed_node_IP_address</i>
Management network netmask	Enter the subnet mask for the appliance.
DNS server list	Enter the DNS servers for the appliance.
NTP servers list	Enter the NTP services for the appliance.
Enable SSH	Deselected
Allow root SSH logins	Selected

13. On the **Ready to complete** page, review the deployment details and click **Finish**.
The NSX Manager virtual machine begins to deploy.

Join the New NSX Manager Node to the NSX Manager Cluster

You join the newly deployed NSX Manager node to the cluster by using the virtual machine web console from the vSphere Client.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu** › **VMs and Templates**.
3. In the inventory expand **vCenter Server** › **Datacenter** › **NSX Folder**.
4. Click the VM of an operational NSX Manager node in the cluster, click **Launch web console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

5. Retrieve the ID of the NSX Manager cluster.

- a) Run the command to view the cluster ID.

```
get cluster config | find Id:
```

- b) Record the cluster ID.

6. Retrieve the API thumbprint of the NSX Manager API certificate.

- a) Run the command to view the certificate API thumbprint.

```
get certificate api thumbprint
```

- b) Record the certificate API thumbprint.

7. Exit the VM Web console.

8. In the vSphere Client, click the VM of the newly deployed NSX Manager node, click **Launch Web console**, and log in by using administrator credentials.

Setting	Value
User name	admin
Password	<i>nsx_admin_password</i>

9. Run the command to join the new NSX Manager node to the cluster.

```
join existing_node_ip cluster-id cluster_id thumbprint api_thumbprint username admin
```

The new NSX Manager node joins the cluster.

Validate the Status of the NSX Manager Cluster

After you added the new NSX Manager node to the cluster, you must validate the operational state of the NSX Manager cluster.

To view the state of the NSX Manager cluster, you log in to the NSX Manager for the particular domain.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Appliances**.
4. Verify that the **Cluster** status is green and *Stable* and that each cluster node is *Available*.

Restore the NSX Manager SSL Certificate

After you add the new NSX Manager node to the cluster and validate the cluster status, you must restore the SSL certificate to the new node.

To view the certificate of the failed NSX Manager cluster node, you log in to the NSX Manager for the domain.

Table 239: NSX Manager Clusters in the SDDC

NSX Manager Cluster	NSX Manager URL
Management domain NSX Manager cluster	<a href="https://<FQDN of management domain NSX Manager>/login.jsp?local=true">https://<FQDN of management domain NSX Manager>/login.jsp?local=true

Table continued on next page

Continued from previous page

NSX Manager Cluster	NSX Manager URL
Workload domain NSX Manager cluster	<code>https://<FQDN of workload domain NSX Manager>/login.jsp?local=true</code>

This procedure is an example for restoring the certificate of a management domain NSX Manager cluster node.

1. In a Web browser, log in to the NSX Manager cluster for the management domain.

Setting	Value
URL	<code>https://<FQDN of management domain NSX Manager>/login.jsp?local=true</code>
User name	admin
Password	<code>nsx_admin_password</code>

2. On the main navigation bar, click **System**.
3. In the left pane, under **Settings**, click **Certificates**.
4. Locate and copy the ID of the certificate that was issued by CA to the node that you are restoring.
5. Run the command to install the CA-signed certificate on the new NSX Manager node.

```
curl -H 'Accept: application/json' -H 'Content-Type: application/json' \
--insecure -u 'admin:nsx_admin_password' -X POST \
'https://nsx_host_node/api/v1/node/services/http
action=apply_certificate&certificate_id=certificate_id'
```

IMPORTANT

If assigning the certificate fails because the certificate revocation list (CRL) verification fails, see <https://kb.vmware.com/kb/78794>. If you disable the CRL checking to assign the certificate, after assigning the certificate, you must re-enable the CRL checking.

Restart the NSX Manager Node

After assigning the certificate, you must restart the new NSX Manager node.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (`https://<vcenter_server_fqdn>/ui`).
2. Select **Menu** > **VMs and Templates**.
3. In the inventory expand **vCenter Server** > **Datacenter** > **NSX Folder**.
4. Right click the new NSX Manager VM and select **Guest OS** > **Restart**.

Validate the Status of the NSX Manager Cluster

After restoring an NSX Manager node, you must validate the system status of the NSX Manager cluster.

To view the system status of the NSX Manager cluster, you log in to the NSX Manager for the particular domain.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (`https://<nsx_manager_cluster_fqdn>/login.jsp?local=true`)
2. On the **Home** page, click **Monitoring Dashboards** > **System**.
3. Verify that all components are healthy.

4. If the host transport nodes are in a `Pending` state, run **Configure NSX** on these nodes to refresh the UI.

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Update or Recreate the VM Anti-Affinity Rule for the NSX Manager Cluster Nodes

During the NSX Manager bring-up process, SDDC Manager creates a VM anti-affinity rule to prevent the VMs of the NSX Manager cluster from running on the same ESXi host. If you redeployed all NSX Manager cluster nodes, you must recreate this rule. If you redeployed one or two nodes of the cluster, you must add the new VMs to the existing rule.

1. In a web browser, log in to the management domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > Hosts and Clusters**.
3. In the inventory expand **vCenter Server > Datacenter**.
4. Click the cluster object.
5. Click the **Configure** tab and click **VM/Host Rules**.
6. Update or recreate the VM anti-affinity rule.
 - If you redeployed one or two nodes of the cluster, add the new VMs to the existing rule.
 1. Click the VM anti-affinity rule name and click **Edit**.
 2. Click **Add VM/Host rule member**, select the new NSX Manager cluster nodes, and click **Add**.
 - If you redeployed all NSX Manager cluster nodes, click **Add VM/Host rule**, enter these values to create the rule, and click **OK**.

Setting	Value
Name	Enter the name of the anti-affinity rule
Type	Separate virtual machines
Members	Click Add VM/Host rule member , select the NSX Manager cluster nodes, and click Add .

Validate the SDDC Manager Inventory State

After a successful restore of an NSX Manager cluster, you must verify that the SDDC Manager inventory is consistent with the recovered virtual machines. You run this verification by using the `sos` tool.

1. Log in to SDDC Manager by using a Secure Shell (SSH).
2. Verify the SDDC Manager health.
 - a) Run the command to view the details about the VMware Cloud Foundation system.

```
sudo /opt/vmware/sddc-support/sos --get-vcf-summary
```

- b) When prompted, enter the `vcf_password`.

All tests show green state.

3. Run the command to collect the log files from the restore of the NSX Manager cluster.

```
sudo /opt/vmware/sddc-support/sos --domain-name domain_name --nsx-logs
```

Refresh the SSH keys that are stored in the SDDC Manager inventory. See [VMware Cloud Foundation SDDC Manager Recovery Scripts \(79004\)](#).

Restoring NSX Edge Cluster Nodes

If one or both NSX Edge cluster nodes fail due to a hardware or software issue, you must redeploy the failed NSX Edge instances. You do not restore the NSX Edge nodes from a backup.

Prepare for Restoring NSX Edge Cluster Nodes

Before restoring an NSX Edge node, you must retrieve its deployment details from the NSX Manager cluster and retrieve the credentials of the failed NSX Edge node from SDDC Manager.

Retrieve the NSX Edge Node Deployment Details from NSX Manager Cluster

Before restoring a failed NSX Edge node, you must retrieve its deployment details from the NSX Manager cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge Transport Nodes** tab.
5. Select the check-box for the failed NSX Edge node.
6. Click **Actions** and select **Change node settings**.
7. Record the **Host name/FQDN** value and click **Cancel**.
8. Click **Actions** and select **Change Edge VM Resource Reservations**.
9. Record the **Existing form factor** value and click **Cancel**.
10. Click the name of the NSX Edge node that you plan to replace and record the following values.
 - Name
 - Management IP
 - Transport Zones
 - Edge Cluster
11. Click **Edit**, record the following values, and click **Cancel**.
 - Edge Switch Name
 - Uplink Profile
 - IP Assignment
 - Teaming Policy Uplink Mapping

Retrieve the NSX Edge Node Credentials from SDDC Manager

Before restoring the failed NSX Edge node that is deployed by SDDC Manager, you must retrieve its credentials from the SDDC Manager inventory.

1. In the SDDC Manager user interface, from the navigation pane click **Developer center**.
2. Click the **API explorer** tab.
3. Expand **APIs for managing credentials** and click **GET /v1/credentials**.
4. In the **resourceName** text box, enter the FQDN of the failed NSX Edge node, and click **Execute**.
5. Under **Response**, click **PageOfCredential** and click each credential ID.
6. Record the user names and passwords for these credentials.

Credential Type	Username	Password
SSH	root	<i>edge_root_password</i>
API	admin	<i>edge_admin_password</i>
AUDIT	audit	<i>edge_audit_password</i>

Retrieve the Workload Domain vSphere Cluster ID from SDDC Manager

If you are restoring a failed workload domain NSX Edge node, you must retrieve the ID of the vSphere cluster for the workload domain. During the restore process, you use this vSphere cluster ID to recreate the vSphere DRS rule name with its original name.

You use the SDDC Manager user interface to retrieve the ID of the vSphere cluster for the workload domain.

1. In the SDDC Manager user interface, from the navigation pane click **Developer center**.
2. Click the **API explorer** tab.
3. Expand **APIs for managing clusters**, click **GET /v1/clusters**, and click **Execute**.
4. Under **Response**, click **PageOfClusters** and click **Cluster**.
5. Record the **ID of the cluster** for the workload domain cluster ID.

Replace the Failed NSX Edge Node with a Temporary NSX Edge Node

You deploy a temporary NSX Edge node in the domain, add it to the NSX Edge cluster, and then delete the failed NSX Edge node.

Deploy a Temporary NSX Edge Node

To avoid conflicts with the failed NSX Edge node, you deploy a temporary NSX Edge node with a new FQDN and IP address.

Allocate the FQDN and IP address for the temporary NSX Edge node for the domain of the failed node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Click **Add edge VM**.
6. On the **Name and description** page, enter these values and click **Next**.

Setting	Value
Name	Enter the VM name
Host name/FQDN	Enter the FQDN
Form factor	Medium

7. On the **Credentials** page, enter these values and the passwords recorded in the earlier steps and then click **Next**.

Setting	Value
CLI user name	admin
CLI password	<i>edge_admin_password</i>
CLI confirm password	<i>edge_admin_password</i>
Allow SSH login	Yes
System root password	<i>edge_root_password</i>
System root password confirm	<i>edge_root_password</i>
Allow root SSH login	No
Audit user name	audit
Audit password	<i>edge_audit_password</i>
Audit confirm password	<i>edge_audit_password</i>

8. On the **Configure deployment** page, select the following and click **Next**.

Setting	Value
Compute manager	Enter the vCenter Server FQDN
Cluster	Select the cluster
Datastore	Select the vSAN datastore

9. On the **Configure node settings** page, enter these values and click **Next**.

Setting	Value
IP Assignment	Static
Management IP	Enter the management IP address.
Default Gateway	Enter the default gateway
Management interface	Select the management network distributed port group
Search domain names	Enter the search domain
DNS servers	Enter the DNS servers
NTP Servers	Enter the NTP servers

10. On the **Configure NSX** page, enter these values which are already recorded and click **Finish**.

Setting	Value
Edge switch name	Enter the edge switch name.
Transport zone	Enter the transport zone names.
Uplink profile	Enter the uplink profile name.
IP assignment	Use static IP list
Static IP list	Enter the static IP list.
Gateway	Enter the gateway IP
Subnet mask	Enter the subnet mask

Table continued on next page

Continued from previous page

Setting	Value
Teaming policy switch mapping	Enter the values for Uplink1 and Uplink2.

Replace the Failed NSX Edge Node with the Temporary NSX Edge Node

You add the temporary NSX Edge node to the NSX Edge cluster by replacing the failed NSX Edge node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge clusters** tab.
5. Select the check-box for the NSX Edge cluster.
6. Click **Action** and select **Replace edge cluster member**.
7. From the **Replace** drop down menu, select the Failed edge node and from the **with** drop down menu, select the Temporary edge node and then click **Save**.

Delete the Failed NSX Edge Node from the NSX Manager Cluster

After replacing the failed NSX Edge node with the temporary NSX Edge node in the NSX Edge cluster, you delete the failed node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Select the check-box for the failed NSX Edge node and click **Delete**.
6. In the confirmation dialog box, click **Delete**.

Validate the Temporary State of the NSX Edge Cluster Nodes

After replacing the failed NSX Edge node with a temporary NSX Edge node, you must verify the state of the NSX Edge cluster nodes.

You validate the state of the temporary NSX Edge node and the second NSX Edge node in the cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Verify all edge transport nodes show these values.

Setting	Value
Configuration state	Success

Table continued on next page

Continued from previous page

Setting	Value
Node status	Up
Tunnels	Upward arrow mark with number of tunnels

Replace the Temporary NSX Edge Node with the Redeployed NSX Edge Node

After you replaced and deleted the failed NSX Edge node, to return the NSX Edge cluster to its original state, you redeploy the failed node, add it to the NSX Edge cluster, and delete then temporary NSX Edge node.

Redeploy the Failed NSX Edge Node

You deploy a new NSX Edge node by using the configurations of the failed NSX Edge node that you retrieved during the preparation for the restore.

To return the NSX Edge cluster to the original state, you must use the FQDN and IP address of the failed NSX Edge node that you deleted. This procedure ensures that the inventory in SDDC Manager is accurate.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.
5. Click **Add edge VM**.
6. On the **Name and description** page, enter these values and click **Next**.

Setting	Value
Name	Enter the VM name
Host name/FQDN	Enter the FQDN
Form factor	Medium

7. On the **Credentials** page, enter these values which are recorded earlier and click **Next**.

Setting	Value
CLI user name	admin
CLI password	<i>edge_admin_password</i>
CLI confirm password	<i>edge_admin_password</i>
Allow SSH login	Yes
System root password	<i>edge_root_password</i>
System root password confirm	<i>edge_root_password</i>
Allow root SSH login	No
Audit user name	audit
Audit password	<i>edge_audit_password</i>
Audit confirm password	<i>edge_audit_password</i>

8. On the **Configure deployment** page, select these values and click **Next**.

Setting	Value
Compute manager	Enter the vCenter Server FQDN
Cluster	Enter the cluster name
Resource pool	Enter the resource pool
Datastore	Enter the datastore

9. On the **Configure Node Settings** page, enter these values and click **Next**.

Setting	Value
IP assignment	Static
Management IP	Enter the management IP address.
Default gateway	Enter the default gateway
Management interface	Select the management network distributed port group
Search domain names	Enter the search domain
DNS servers	Enter the DNS servers
NTP servers	Enter the NTP servers

10. On the **Configure NSX** page, enter these values which are recorded earlier and click **Finish**.

Setting	Value
Edge switch name	Enter the edge switch name.
Transport zone	Enter the transport zone names.
Uplink profile	Enter the uplink profile name.
IP assignment	Use static IP list
Static IP list	Enter the static IP list.
Gateway	Enter the gateway IP
Subnet mask	Enter the subnet mask
Teaming policy switch mapping	Enter the values for Uplink1 and Uplink2.

Replace the Temporary NSX Edge Node with the Redeployed NSX Edge Node

After deploying the new NSX Edge node with the same configuration as the failed NSX Edge node, you replace the temporary NSX Edge node with the redeployed failed node in the NSX- Edge cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge clusters** tab.
5. Select the check-box for the NSX Edge cluster.
6. Click **Action** and select **Replace edge cluster member**.
7. From the **Replace** drop down menu, select the temporary node and from the **with** drop down menu, select the new node and then click **Save**.

Delete the Temporary NSX Edge Node

After replacing the temporary NSX Edge node with the new NSX Edge node in the NSX Edge cluster, you delete the temporary node.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes > .**
4. Click the **Edge transport nodes** tab.
5. Select the check-box for the temporary NSX Edge node and click **Delete**.
6. In the confirmation dialog box, click **Delete**.

Update or Recreate the VM Anti-Affinity Rule for the NSX Edge Cluster Nodes

During the NSX Edge deployment process, SDDC Manager creates a VM anti-affinity rule to prevent the nodes of the NSX Edge cluster from running on the same ESXi host. If you redeployed the two NSX Edge cluster nodes, you must recreate this rule. If you redeployed one node of the cluster, you must add the new VM to the existing rule.

1. In a web browser, log in to the domain vCenter Server by using the vSphere Client (https://<vcenter_server_fqdn>/ui).
2. Select **Menu > Hosts and Clusters**.
3. In the inventory expand **vCenter Server > Datacenter**.
4. Click the cluster object.
5. Click the **Configure** tab and click **VM/Host Rules**.
6. Update or recreate the VM anti-affinity rule.
 - If you redeployed one of the nodes in the NSX Edge cluster, add the new VM to the existing rule.
 1. Click the VM anti-affinity rule name and click **Edit**.
 2. Click **Add VM/Host rule member**, select the new NSX Edge cluster node, and click **Add**.
 - If you redeployed the two nodes in the NSX Edge cluster, click **Add VM/Host rule**, enter these values to create the rule, and click **OK**.

Setting	Value
Name	Enter the name of the anti-affinity rule
Type	Separate virtual machines
Members	Click Add VM/Host rule member , select the NSX Edge cluster nodes, and click Add .

Validate the State of the NSX Edge Cluster Nodes

After replacing the temporary NSX Edge node with the redeployed failed NSX Edge node, you must verify the state of the NSX Edge cluster nodes.

You validate the state of the redeployed NSX Edge node and the second NSX Edge node in the cluster.

1. In a web browser, log in to the NSX Manager cluster for the domain by using the user interface (https://<nsx_manager_cluster_fqdn>/login.jsp?local=true)
2. On the main navigation bar, click **System**.
3. In the left pane, under **Configuration**, click **Fabric > Nodes**.
4. Click the **Edge transport nodes** tab.

5. Verify all edge transport nodes show these values.

Setting	Value
Configuration state	Success
Node status	Up
Tunnels	Upward arrow mark with number of tunnels

Image-Based Backup and Restore of VMware Cloud Foundation

For an image-based backup of the VMware Cloud Foundation, use a solution compatible with the VMware vSphere Storage APIs - Data Protection (formerly known as VMware vStorage APIs for Data Protection or VADP).

vSphere Storage APIs - Data Protection compatible backup software connects to the vCenter servers in the management domain to perform backups. In the event of failure, the backup software connects to the vCenter servers in the management domain to restore the VMs. If the management domain is lost, the vCenter servers are no longer available and must be restored first. Choosing a backup software that supports Direct Restore to an ESXi host allows restoring the vCenter Servers.

Connect your backup solution with the management domain vCenter Server and configure it. To reduce the backup time and storage cost, use incremental backups in addition to the full ones.

Acquiesced backups are enabled for VMware Aria Suite Lifecycle and Workspace ONE Access.

Upgrading to VMware Cloud Foundation 5.2.x on Dell VxRail

The following procedures provide information about upgrading to VMware Cloud Foundation 5.2.x on Dell VxRail.

NOTE

Review the [VMware Interoperability Matrix](#) to verify compatibility and upgradability before planning and starting an upgrade.

You can perform a sequential or skip-level upgrade to VMware Cloud Foundation 5.2.x on Dell VxRail from VMware Cloud Foundation 4.5 or later. If your environment is at a version earlier than 4.5, you must upgrade the management domain and all VI workload domains to VMware Cloud Foundation 4.5 or later and then upgrade to VMware Cloud Foundation 5.2.x.

WARNING

vSphere with Tanzu enabled clusters, may require a specific upgrade sequence. See [KB 88962](#) for more information.

The first step is to download the bundles for each VMware Cloud Foundation on Dell VxRail component that requires an upgrade. After all of the bundles are available in SDDC Manager, upgrade the management domain and then your VI workload domains.

- [Downloading VMware Cloud Foundation Upgrade Bundles](#)
- [Upgrade the Management Domain to VMware Cloud Foundation 5.2.x](#)
- [Upgrade VI Workload Domains to VMware Cloud Foundation 5.2.x](#)

SDDC Manager Functionality During an Upgrade to VMware Cloud Foundation 5.2

During the upgrade to VMware Cloud Foundation 5.2, some SDDC Manager functionality may be limited during each phase of the upgrade. Prior to initiating the upgrade determine if you will need to perform any of these tasks.

Upgrade States and Terminology

- **Source BOM** - Prior to initiating the upgrade all components are at VMware Cloud Foundation 4.5.x, 5.0, or 5.1.
- **SDDC Manager only** - You have updated SDDC Manager to 5.2, but none of the other BOM components.
- **Split BOM** - Management domain or VI Workload Domain is only partially updated to VMware Cloud Foundation 5.2.
- **Mixed 4.5.x/5.x BOM** - Some workload domains (Management or VI) have been completely upgraded to VMware Cloud Foundation 5.2 and at least one VI Workload Domain is at the Source 4.5.x BOM version.
- **Mixed 5.x BOM** - Some workload domains (Management or VI) have been completely upgraded to VMware Cloud Foundation 5.2 and at least one VI Workload Domain is at the Source 5.0 or 5.1 BOM version.
- **Target BOM** - All components are at VMware Cloud Foundation 5.2.

When a VMware Cloud Foundation instance is in Source BOM or Target BOM, the features available within SDDC Manager are as expected for that given release. However when in a Mixed BOM the operations available vary per workload domain depending on which state the domain itself is in.

The following table indicates the functions available within SDDC Manager during an upgrade.

Table 240: SDDC Manager Functionality During Upgrade

Category	Feature	SDDC Manager only	Split BOM	Mixed 4.5.x/5.x BOM	Mixed 5.x BOM
Backup / Restore	Configure and perform Backup / Restore	Y	Y	Y	Y
CEIP	Activate / Deactivate CEIP	Y	Y	Y	Y
Certificate Management	View/Generate/ Upload/Install	Y	Y	Y	Y
NSX Edge Cluster	Expand edge cluster	Y	Y	Y	Y
DNS / NTP configuration	Validate / Configure DNS	Y	Y	Y	Y
	Validate / Configure NTP	Y	Y	Y	Y
Licensing	Update License Key Information	Y	Y	Y	Y
	Add License Key	Y	Y	Y	Y
	Relicensing	Y	Y	Y	Y
	License check	Y	Y	Y	Y
LCM	Connect to VMware or Dell Depot / Download Bundles	Y	Y	Y	Y
	LCM Pre checks	Y	Y	Y	Y
	Schedule Bundle Download	Y	Y	Y	Y
	Install vCenter Patch	Y	Y	Y	Y
	Install ESXi Patch	Y	Y	Y	Y
	Install NSX Patch	Y	Y	Y	Y

Table continued on next page

Continued from previous page

Category	Feature	SDDC Manager only	Split BOM	Mixed 4.5.x/5.x BOM	Mixed 5.x BOM
Password Management	Rotate/Update/Retry/Cancel	Y	Y	Y	Y
User Operations	Add / Remove User / Group	Y	Y	Y	Y
Workload Domain	Add/Remove ESXi Host	Y	Y	Y	Y
	Add/Remove vSphere Cluster	Y	Y	Y	Y
	Add 4.5.x Workload Domain	Y If the management domain is at 4.5.x. NOTE Contact Broadcom Support for a workload and if the management domain is at 5.x.	Y If the management domain is at 4.5.x. NOTE Contact Broadcom Support for a workload and if the management domain is at 5.x.	Y If the management domain is at 4.5.x. NOTE Contact Broadcom Support for a workload and if the management domain is at 5.x.	N/A
	Add 5.x Workload Domain in ELM mode	Y	Y	Y	Y
	Add 5.x Isolated Workload Domain	Y	Y	Y	Y
	Remove 4.5.x Workload Domain	Y	Y	Y	N/A
	Remove 5.0 Workload Domain	Y	Y	Y	Y
	Stretch a vSphere Cluster	You cannot stretch clusters in 4.5.x workload domains, but can stretch cluster in 5.x workload domains.	You cannot stretch clusters in 4.5.x workload domains, but can stretch cluster in 5.x workload domains.	You cannot stretch clusters in 4.5.x workload domains, but can stretch cluster in 5.x workload domains.	Y
Expand a Stretched vSphere Cluster	You cannot expand clusters in 4.5.x workload domains, but can expand clusters in	You cannot expand clusters in 4.5.x workload domains, but can expand clusters in	You cannot expand clusters in 4.5.x workload domains, but can expand clusters in	Y	

Table continued on next page

Continued from previous page

Category	Feature	SDDC Manager only	Split BOM	Mixed 4.5.x/5.x BOM	Mixed 5.x BOM
		5.x workload domains.	5.x workload domains.	5.x workload domains.	
	Shrink a Stretched vSphere Cluster	You cannot shrink clusters in 4.5.x workload domains, but can shrink clusters in 5.x workload domains.	You cannot shrink clusters in 4.5.x workload domains, but can shrink clusters in 5.x workload domains.	You cannot shrink clusters in 4.5.x workload domains, but can shrink clusters in 5.x workload domains.	Y

vSphere UI Client Plug-ins

Identify all vSphere UI client plug-ins prior to the upgrade.

It may be possible to upgrade some vSphere UI client plug-ins before upgrading to vSphere 8.0. Contact your 3rd Party vendor to determine the best upgrade path.

Monitor VMware Cloud Foundation Updates

You can monitor in-progress updates for VMware Cloud Foundation components.

1. In the In-Progress Updates section, click **View Status** to view the high-level update progress and the number of components to be updated.
2. Details of the component being updated is shown below that. The image below is an example and may not reflect the actual versions.

VMware Cloud Foundation Update Status

VMware Cloud Foundation Update 4.0.1.0

Released 06/23/2020 11 GB

This VMware Cloud Foundation Upgrade 4.0.0.1 to 4.0.1.0 contains features, critical bugs and security fixes
<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.01/rn/VMware-Cloud-Foundation-401-Release-N>

Updating COMMON SERVICES

> SDDC MANAGER

 In Progress

3. Click the arrow to see a list of tasks being performed to update the component. As the task is completed, it shows a green check mark.

VMware Cloud Foundation Update Status

VMware Cloud Foundation Update 4.0.1.0

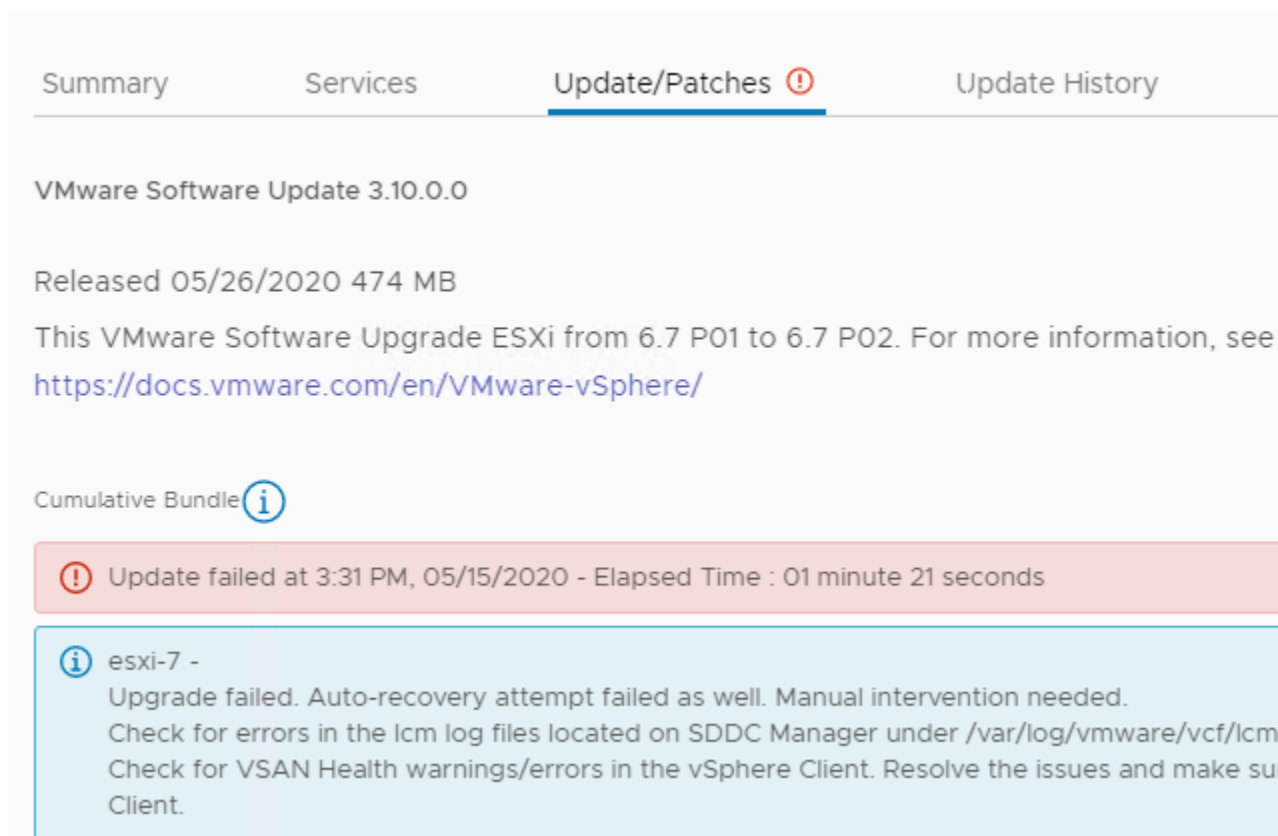
Released 06/23/2020 11 GB

This VMware Cloud Foundation Upgrade 4.0.0.1 to 4.0.1.0 contains features, critical bugs and security f
<https://docs.vmware.com/en/VMware-Cloud-Foundation/4.01/rn/VMware-Cloud-Foundation-401-Release>

Updating OPERATIONS MANAGER

<div style="display: flex; align-items: center;"> <div style="margin-right: 10px;"> } </div> <div> <p>SDDC MANAGER</p> <p>> COMMON SERVICES</p> <p>> OPERATIONS MANAGER</p> <p>DOMAIN MANAGER</p> <p>SOLUTIONS MANAGER</p> <p>SDDC MANAGER UI</p> <p>LCM</p> <p>MULTI SITE SERVICE</p> </div> </div>	<div style="display: flex; align-items: center; margin-bottom: 10px;"> ○ In Progress </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> ✓ Updated 🕒 </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> ○ In Progress </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> 🕒 Queued </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> 🕒 Queued </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> 🕒 Queued </div> <div style="display: flex; align-items: center; margin-bottom: 10px;"> 🕒 Queued </div> <div style="display: flex; align-items: center;"> 🕒 Queued </div>
---	--

4. When all tasks to update a component have been completed, the update status for the component is displayed as Updated.
5. If a component fails to be updated, the status is displayed as Failed. The reason for the failure as well as remediation steps are displayed. The image below is an example and may not reflect the actual versions in your environment.



Summary Services **Update/Patches** Update History

VMware Software Update 3.10.0.0

Released 05/26/2020 474 MB

This VMware Software Upgrade ESXi from 6.7 P01 to 6.7 P02. For more information, see <https://docs.vmware.com/en/VMware-vSphere/>

Cumulative Bundle

! Update failed at 3:31 PM, 05/15/2020 - Elapsed Time : 01 minute 21 seconds

i esxi-7 -
 Upgrade failed. Auto-recovery attempt failed as well. Manual intervention needed.
 Check for errors in the lcm log files located on SDDC Manager under /var/log/vmware/vcf/lcm.
 Check for VSAN Health warnings/errors in the vSphere Client. Resolve the issues and make sure Client.

6. After you resolve the issues, you can retry the update.

View VMware Cloud Foundation Update History

The Update History page displays all updates applied to a workload domain.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. Click the name of a workload domain and then click the **Update History** tab.
 All updates applied to this workload domain are displayed. If an update bundle was applied more than once, click **View Past Attempts** to see more information.

Access VMware Cloud Foundation Upgrade Log Files

You can check the log files for failed upgrades to help troubleshoot and resolve issues.

1. SSH in to the SDDC Manager appliance with the `vcf` user name and enter the password.
2. To access upgrade logs, navigate to the `/var/log/vmware/vcf/lcm` directory.
 - `lcm-debug` log file contains debug level logging information.
 - `lcm.log` contains information level logging.
3. To create an sos bundle for support, see Supportability and Serviceability (SoS) Utility in the *VMware Cloud Foundation Administration Guide*.

Downloading VMware Cloud Foundation Upgrade Bundles

Before you can upgrade VMware Cloud Foundation, you must download the upgrade bundles for each VMware Cloud Foundation component that requires an upgrade.

Online and Offline Downloads

If the SDDC Manager appliance can connect to the internet (either directly or through a proxy server), you can download upgrade bundles from the VMware Depot and the Dell Depot.

If the SDDC Manager appliance cannot connect to the internet, you can use the Bundle Transfer Utility or connect to an offline depot.

See [Public URL list for SDDC Manager](#) for information about the URLs that must be accessible to download bundles.

Other Bundle Types

In addition to upgrade bundles, VMware Cloud Foundation includes the following bundle types:

- **Install Bundles**
An install bundle includes software binaries to install VI workload domains (vCenter Server and NSX) and VMware Aria Suite Lifecycle. You download install bundles using the same process that you use for upgrade bundles.
- **Async Patch Bundles**
An async patch bundle allows you to apply critical patches to certain VMware Cloud Foundation components (NSX Manager, and vCenter Server) when an update or upgrade bundle is not available. If you are running VMware Cloud Foundation 5.1 or earlier, you must use the Async Patch Tool to download an async patch bundle. See [Async Patch Tool](#). Starting with VMware Cloud Foundation 5.2, you can download async patches using the SDDC Manager UI or Bundle Transfer Utility.

Connect SDDC Manager to a Software Depot for Downloading Bundles

SDDC Manager can connect to a software depot to download software bundles, compatibility data, and more.

To connect to the online depot, SDDC Manager must be able to connect to the internet, either directly or through a proxy server.

To connect to an offline depot, you must first configure it. See [KB 312168](#) for information about the requirements and process for creating an offline depot. To download bundles to an offline depot, see "Download Bundles to an Offline Depot" in the *VMware Cloud Foundation Lifecycle Management Guide*.

SDDC Manager supports two types of software depots:


- Online depot
- Offline depot

You can only connect SDDC Manager to one type of depot. If SDDC Manager is connected to an online depot and you configure a connection to an offline depot, the online depot connection is disabled and deleted.

1. In the navigation pane, click **Administration** > **Depot Settings**.

Online Depots

VMware Depot

 Depot connection not set up. Authenticate your Customer Connect Account to set it up.

AUTHENTICATE

Offline Depot

Offline Depot

 Depot connection not set up.

SET UP

2. Connect SDDC Manager to an online depot or an offline depot.

Depot Type	Configuration Steps
Online	<ol style="list-style-type: none"> 1. Click Authenticate for the VMware Depot. 2. Type your Broadcom Support Portal user name and password. 3. Type your Dell Depot user name and password. 4. Click Authenticate
Offline	<ol style="list-style-type: none"> 1. Click Set Up for the Offline Depot. 2. Enter the following information for the offline depot: <ul style="list-style-type: none"> – FQDN or IP address – Port number – User name – Password 3. Click Set Up.

SDDC Manager attempts to connect to the depot. If the connection is successful, SDDC Manager starts looking for available bundles. To view available bundles, click **Lifecycle Management > Bundle Management** and then click the **Bundles** tab. It may take some time for all available bundles to appear.

Download Bundles Using SDDC Manager

After you connect SDDC Manager to an online or offline depot, you can view and download available upgrade bundles.

Connect SDDC Manager to an online or offline depot. See [Connect SDDC Manager to a Software Depot for Downloading Bundles](#).

If SDDC Manager does not have direct internet access, configure a proxy server or use the Bundle Transfer Utility for offline bundle downloads.

- [Configure a Proxy Server for Downloading Bundles](#)
- [Offline Bundle Download for VMware Cloud Foundation](#)

When you download bundles, SDDC Manager verifies that the file size and checksum of the downloaded bundles match the expected values.

1. In the navigation pane, click **Lifecycle Management** › **Bundle Management**.
2. Click the **Bundles** tab to view available bundles.

NOTE

If you just connected SDDC Manager to a depot, it can take some time for bundles to appear.

All available bundles are displayed. Install bundles display an Install Only Bundle label. If the bundle can be applied right away, the Bundle Details column displays the workload domains to which the bundle needs to be applied to, and the Availability column says Available. If another bundle must be applied before a particular bundle, the Availability field displays Future.

To view more information about the bundle, click **View Details**. The Bundle Details section displays the bundle version, release date, and additional details about the bundle.

3. For the bundle you want to download, do one of the following:

- Click **Download Now** for an immediate download.

The bundle download begins right away.

- Click **Schedule Download** to schedule a download.

Select the date and time for the bundle download and click **Schedule**.

4. Click the **Download History** tab to see the downloaded bundles.

Configure a Proxy Server for Downloading VMware Cloud Foundation Bundles

If SDDC Manager does not have direct internet access, you can configure a proxy server to download bundles. VMware Cloud Foundation 5.2 and later support proxy servers with authentication.

1. In the navigation pane, click **Administration** › **Proxy Settings**.
2. Click **Set Up Proxy**.
3. Toggle the **Enable Proxy** setting to the on position.
4. Select **HTTP** or **HTTPS**.
5. Enter the proxy server IP address and port number.
6. If your proxy server requires authentication, toggle the **Authentication** setting to the on position and enter the user name and password.
7. Click **Save**.

You can now download bundles as described in [Download Bundles Using SDDC Manager](#).

Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles

If the SDDC Manager appliance does not have access to the VMware Depot and the Dell Depot, you can use the Bundle Transfer Utility to download the bundles to a different computer and then upload them to the SDDC Manager appliance.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.

NOTE

The Bundle Transfer Utility is the only supported method for downloading bundles. Do not use third-party tools or other methods to download bundles.

Using the Bundle Transfer Utility to upgrade to VMware Cloud Foundation 5.2.x involves the following steps:

- Download the latest version of the Bundle Transfer Utility.
- On a computer with access to the internet, use the Bundle Transfer Utility to download the bundles and other required files.
- Copy the bundles and other required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and other required files to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS. To use this option, the proxy certificate must be added to Bundle Transfer Utility JRE default trust store. For example: <pre>/opt/obtu/jre/lin64/bin/keytool -importcert -file ca-bundle.crt -keystore /opt/obtu/jre/lin64/lib/ security/cacerts</pre>
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUserUsername--proxyServer
proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../password.txt --
proxyHttps
```

1. Download the most recent version of the Bundle Transfer Utility on a computer with internet access.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
 - b) Click the version of VMware Cloud Foundation to which you are upgrading.
 - c) Click **Drivers & Tools**.
 - d) Click the download icon for the Bundle Transfer Utility.
 - e) Extract `lcm-tools-prod.tar.gz`.
 - f) Navigate to the `lcm-tools-prod/bin/` and confirm that you have execute permission on all folders.

2. Download bundles and other artifacts to the computer with internet access.

- a) Download the manifest file.

This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUserUsername
```

For `--depotUser`, enter your Broadcom Support Portal user name.

Note the location to which the Bundle Transfer Utility downloads the manifest. You will use this as the `--sourceManifestDirectory` when you upload the manifest. For example:

```
Validating the depot user credentials...
Downloading LCM Manifest to: /root/PROD2/evo/vmw
Successfully completed downloading file
```

- b) Download the compatibility data.

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUserUsername-
pdudell_depot_email
```

To specify a download location, use `--outputDirectory` followed by the path to the directory.

- c) Download the vSAN HCL file.

```
./lcm-bundle-transfer-util --vsanHclDownload
```

- d) Download the upgrade bundles.

```
./lcm-bundle-transfer-util --download "downloadPartnerBundle" --download
"withCompatibilitySets" --outputDirectoryabsolute-path-output-dir--depotUsercus
tomer_connect_email--svcurrent-vcf-version--ptarget-vcf-version--pdudell_depot_
email
```

where

<code>absolute-path-output-dir</code>	Path to the directory where the bundle files should be downloaded. This directory folder must have 777 permissions. If you do not specify the download directory, bundles are downloaded to the default directory with 777 permissions.
<code>depotUser</code>	User name for the Broadcom Support Portal. You will be prompted to enter the depot user password. If there are any special characters in the password, specify the password within single quotes.
<code>current-vcf-version</code>	Current version of VMware Cloud Foundation. For example, 4.5.2.0.

Table continued on next page

Continued from previous page

<code>target-vcf-version</code>	Target version of VMware Cloud Foundation. For example, 5.2.1.0.
<code>dell_depot_email</code>	Dell depot email address.

e) Specify the bundles to download.

Enter one of the following options:

- all
- install
- patch

You can also enter a comma-separated list of bundle names to download specific bundles. For example: `bundle-38371, bundle-38378`.

Download progress for each bundle is displayed. Wait until all bundles are downloaded successfully.

3. Copy the following files/directories to the SDDC Manager appliance.

- Bundle Transfer Utility
- Manifest file
- Compatibility data files (`VmwareCompatibilityData.json` and `VxrailCompatibilityData.json`)
- vSAN HCL
- Entire bundle output directory

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

NOTE

Make sure to copy the entire output directory, including any VxRail bundles and JSON files.

4. If you downloaded VxRail bundles:

- a) Copy the partner bundle to the `/nfs/vmware/vcf/nfs-mount/bundle/depot/local/bundles` directory on the SDDC Manager appliance.
- b) Copy `partnerBundleMetadata.json` to the `/nfs/vmware/vcf/nfs-mount/bundle/depot/local` directory on the SDDC Manager appliance.
- c) Copy `softwareCompatibilitySets.json` to the `/nfs/vmware/vcf/nfs-mount/bundle/depot/local` directory on the SDDC Manager appliance.
- d) Run following commands on the SDDC Manager appliance:

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/bundle/depot/local
chmod -R 755 /nfs/vmware/vcf/nfs-mount/bundle/depot/local
```

5. Copy the bundle transfer utility to the SDDC Manager appliance.

- a) SSH in to the SDDC Manager appliance using the `vcf` user account.
- b) Enter `su` to switch to the root user.
- c) Create the `lcm-tools` directory.

```
mkdir /opt/vmware/vcf/lcm/lcm-tools
```

NOTE

If the `/opt/vmware/vcf/lcm/lcm-tools` directory already exists with an older version of the Bundle Transfer Utility, you need to delete contents of the existing directory before proceeding.

d) Copy the Bundle Transfer Utility file (`lcm-tools-prod.tar.gz`) that you downloaded in step 1 to the `/opt/vmware/vcf/lcm/lcm-tools` directory.

e) Extract the contents of `lcm-tools-prod.tar.gz`.

```
tar -xvf lcm-tools-prod.tar.gz
```

f) Set the permissions for the `lcm-tools` directory.

```
cd /opt/vmware/vcf/lcm/
chown vcf_lcm:vcf -R lcm-tools
chmod 750 -R lcm-tools
```

6. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and artifacts.

a) Upload the manifest file.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectoryManifest-Directory-
-sddcMgrFqdn FQDN--sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

b) Upload the compatibility files.

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --
inputDirectorycompatibility-file-directory--sddcMgrFqdnFQDN--sddcMgrUserUsernam
e
```

c) Upload the HCL file.

```
./lcm-bundle-transfer-util --vsanHclUpload --inputDirectoryhcl-file-path--
sddcMgrFqdn.sddc-manager-fqdn--sddcMgrUseruser
```

d) Upload the bundle directory.

```
./lcm-bundle-transfer-util --upload "uploadPartnerBundle" --
bundleDirectoryabsolute-path-bundle-dir
```

Offline Download of Independent SDDC Manager Bundles

Once SDDC Manager is upgraded to 5.2 or later, new functionality is introduced that allows you to get the latest SDDC Manager features and security fixes without having to upgrade the entire VMware Cloud Foundation BOM. This procedure describes using the Bundle Transfer Utility to download an SDDC Manager bundle released independently of the VMware Cloud Foundation BOM when SDDC Manager is not connected to an online depot.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- The computer with internet connectivity and the SDDC Manager appliance must have the latest version of the Bundle Transfer Utility installed and configured. See [Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles](#) for more information.

An independent SDDC Manager release includes a fourth digit in its version number, for example SDDC Manager 5.2.0.1.

- On a computer with access to the internet, use the Bundle Transfer Utility to download the independent SDDC Manager bundle and other required files.
- Copy the bundle and other required files to the SDDC Manager appliance.

- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundle and other required files to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download bundles and other artifacts to the computer with internet access.

a) Download the manifest file.

This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username
```

For `--depotUser`, enter your Broadcom Support Portal user name.

b) Download the compatibility data.

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUser
Username --pdu dell_depot_email
```

To specify a download location, use `--outputDirectory` followed by the path to the directory.

c) Download the independent SDDC Manager upgrade bundle.

```
./lcm-bundle-transfer-util --download --sddcMgrVersion four-digit-sddc-version
--depotUser Username --outputDirectory absolute-path-output-dir
```

where

<code>depotUser</code>	User name for the Broadcom Support Portal. You will be prompted to enter the user password. If there are any special characters in the password, specify the password within single quotes.
<code>four-digit-sddc-version</code>	Target version of SDDC Manager. For example, 5.2.0.1.

Table continued on next page

Continued from previous page

<code>absolute-path-output-dir</code>	Path to the directory where the bundle files should be downloaded. This directory folder must have 777 permissions. If you do not specify the download directory, bundles are downloaded to the default directory with 777 permissions.
---------------------------------------	--

Follow the prompts in the Bundle Transfer Utility.

2. Copy the following files/directories to the SDDC Manager appliance.

- Manifest file
- Compatibility data files (`VmwareCompatibilityData.json` and `VxrailCompatibilityData.json`)
- Entire bundle output directory

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

3. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and artifacts.

a) Upload the manifest file.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory Manifest-Directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

b) Upload the compatibility files.

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --inputDirectory compatibility-file-directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

c) Upload the bundle directory.

```
./lcm-bundle-transfer-util --upload --bundleDirectory absolute-path-bundle-dir
```

After the upload completes successfully, you can use the SDDC Manager UI to upgrade SDDC Manager. See [Independent SDDC Manager Upgrade using the SDDC Manager UI](#).

Offline Download of Async Patch Bundles

Once SDDC Manager is upgraded to 5.2 or later, a new option for patching VMware Cloud Foundation components is available in the SDDC Manager UI. This procedure describes using the Bundle Transfer Utility to download async patches when SDDC Manager is not connected to an online depot.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- The computer with internet connectivity and the SDDC Manager appliance must have the latest version of the Bundle Transfer Utility installed and configured. See [Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles](#) for more information.

Offline download of async patches involves the following steps:

- On a computer with access to the internet, use the Bundle Transfer Utility to download the async patch bundle and other required files.
- Copy the bundle and other required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundle and other required files to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download bundles and other artifacts to the computer with internet access.

a) Download the manifest file.

This is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username
```

For `--depotUser`, enter your Broadcom Support Portal user name.

b) Download the compatibility data.

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUser
Username --pdu dell_depot_email
```

To specify a download location, use `--outputDirectory` followed by the path to the directory.

c) Download the product version catalog.

```
./lcm-bundle-transfer-util --depotUser Username --download
productVersionCatalog --outputDirectory directory-path
```

d) List the available async patches.

```
./lcm-bundle-transfer-util --listAsyncPatchBundles listAsyncPatchPartnerBundle
--depotUser Username
```

e) Download an async patch.

```
./lcm-bundle-transfer-util --download --bundle bundle-number --depotUser
Username
```

For example:

```
./lcm-bundle-transfer-util --download --bundle bundle-12345 --depotUser
user@example.com
```

2. Copy the following files/directories to the SDDC Manager appliance.

- Manifest file
- Compatibility data files (`VmwareCompatibilityData.json` and `VxrailCompatibilityData.json`)
- Entire bundle output directory

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

3. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the bundles and artifacts.

a) Upload the manifest file.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory Manifest-Directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

b) Upload the compatibility files.

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --inputDirectory compatibility-file-directory --sddcMgrFqdn FQDN --sddcMgrUser Username
```

c) Upload the product version catalog.

```
./lcm-bundle-transfer-util --upload productVersionCatalog --inputDirectory directory-path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

d) Upload the bundle directory.

```
./lcm-bundle-transfer-util --upload --bundle bundle-number --bundleDirectory absolute-path-bundle-dir
```

- Replace *number* with the bundle number you are uploading. For example: 12345 for `bundle-12345`.
- Replace *absolute-path-bundle-dir* with the path to the location where you copied the output directory. For example: `/nfs/vmware/vcf/nfs-mount/upgrade-bundles`.

After the upload completes successfully, you can use the SDDC Manager UI to apply the async patch. See [Patching the Management and Workload Domains](#).

Offline Download of Flexible BOM Upgrade Bundles

Once SDDC Manager is upgraded to version 5.2 or later, new functionality is introduced to the upgrade planner that allows you to select specific target versions for each VMware Cloud Foundation component you want to upgrade. This procedure describes using the Bundle Transfer Utility to download the bundles for a flexible BOM upgrade when SDDC Manager is not connected to an online depot.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
- A Windows or Linux computer with access to the SDDC Manager appliance for uploading the bundles.
- To upload the manifest file from a Windows computer, you must have OpenSSL installed and configured.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- The computer with internet connectivity and the SDDC Manager appliance must all have the latest version of the Bundle Transfer Utility installed and configured. See [Offline Download of VMware Cloud Foundation 5.2.x Upgrade Bundles](#) for more information.

After you download the bundles, you can use the upgrade planner in the SDDC Manager UI to select any supported version for each of the VMware Cloud Foundation BOM components. This includes async patch versions as well as VCF BOM versions.

Offline download of flexible BOM upgrade bundles involves the following steps:

- On a computer with access to the internet, use the Bundle Transfer Utility to download the required files.
- Copy the required files to the SDDC Manager appliance.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to upload the required files to the internal LCM repository.
- Plan the upgrade using the SDDC Manager UI.
- On the SDDC Manager appliance, use the Bundle Transfer Utility to generate the `plannerFile.json`.
- Copy `plannerFile.json` to the computer with internet access.
- On the computer with access to the internet, download bundles using `plannerFile.json`.
- Copy the bundle directory to the SDDC Manager appliance and use the Bundle Transfer Utility to upload the bundles to the internal LCM repository.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../
password.txt --proxyHttps
```

1. Download the required files to the computer with internet access.

```
./lcm-bundle-transfer-util --download --manifestDownload --depotUser Username --
outputDirectory directory-path
```

The manifest is a structured metadata file that contains information about the VMware product versions included in the release Bill of Materials.

For `--depotUser`, enter your Broadcom Support Portal user name.

```
./lcm-bundle-transfer-util --download --bundleManifests --depotUser Username --
bundleManifestsDir directory-path
```

```
./lcm-bundle-transfer-util --download --compatibilityMatrix --depotUser Username --
pdu dell_depot_email --outputDirectory directory-path
```

```
./lcm-bundle-transfer-util --depotUser Username --download productVersionCatalog --
outputDirectory directory-path
```

```
./lcm-bundle-transfer-util --depotUser Username --download partnerBundleMetadata
```

2. Copy the entire output directory to the SDDC Manager appliance.

You can select any location on the SDDC Manager appliance that has enough free space available. For example, `/nfs/vmware/vcf/nfs-mount/`.

3. On the SDDC Manager appliance, upload/update the files.

```
./lcm-bundle-transfer-util --update --sourceManifestDirectory directory-path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

Use your vSphere SSO credentials for the `--sddcMgrUser` parameter.

```
./lcm-bundle-transfer-util --upload --bundleManifests --bundleManifestsDir directory-path
```

```
./lcm-bundle-transfer-util --update --compatibilityMatrix --inputDirectory directory-path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

```
./lcm-bundle-transfer-util --upload productVersionCatalog --inputDirectory directory-path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

```
./lcm-bundle-transfer-util --upload partnerBundleMetadata --inputDirectory directory-path --sddcMgrFqdn FQDN --sddcMgrUser Username
```

4. In the SDDC Manager UI, plan the upgrade.

See [Flexible BOM Upgrade in VMware Cloud Foundation](#).

5. On the SDDC Manager appliance, use the Bundle Transfer Utility to generate a planner file.

```
./lcm-bundle-transfer-util --generatePlannerFile --sddcMgrUser Username --sddcMgrFqdn FQDN --outputDirectory directory-path --domainNames domain-name -p target-vcf-version
```

For example:

```
./lcm-bundle-transfer-util --generatePlannerFile --sddcMgrUser administrator@vsphere.local --sddcMgrFqdn sddc-manager.example.com --outputDirectory /home/vcd --domainNames mgmt-domain -p 5.2.0.0
```

6. Copy `plannerFile.json` file to the computer with access to the internet.

7. On the computer with access to the internet, download the bundles using the `plannerFile.json`.

```
./lcm-bundle-transfer-util --download downloadPartnerBundle --plannerFile directory-path --depotUser Username --partnerDepotUser dell_depot_email
```

8. Copy the entire output directory to the SDDC Manager appliance.

9. If you downloaded VxRail bundles:

- Copy the partner bundle to the `/nfs/vmware/vcf/nfs-mount/bundle/depot/local/bundles` directory on the SDDC Manager appliance.
- Copy `partnerBundleMetadata.json` to the `/nfs/vmware/vcf/nfs-mount/bundle/depot/local` directory on the SDDC Manager appliance.
- Copy `softwareCompatibilitySets.json` to the `/nfs/vmware/vcf/nfs-mount/bundle/depot/local` directory on the SDDC Manager appliance.
- Run following commands on the SDDC Manager appliance:

```
chown -R vcf_lcm:vcf /nfs/vmware/vcf/nfs-mount/bundle/depot/local
```

```
chmod -R 755 /nfs/vmware/vcf/nfs-mount/bundle/depot/local
```

10. Upload the bundle directory to the SDDC Manager appliance internal LCM repository.

```
./lcm-bundle-transfer-util --upload "uploadPartnerBundle" --bundleDirectory  
directory-path
```

In the SDDC Manager UI browse to the Available Updates screen for the workload domain you are upgrading and click **Schedule Update** or **Update Now** to update the first component. Continue to update the VCF BOM components until they are all updated.

HCL Offline Download for VMware Cloud Foundation

If the SDDC Manager appliance does not have access to the VMware Depot, you can use the Bundle Transfer Utility to manually download the HCL file from the depot on your local computer and then upload it to the SDDC Manager appliance.

- A Windows or Linux computer with internet connectivity (either directly or through a proxy) for downloading the HCL. To upload the HCL file from a Windows computer, you must have OpenSSL installed and configured.
- Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.

NOTE

The Bundle Transfer Utility is the only supported method for downloading HCL. Do not use third-party tools or other methods to download HCL.

If the computer with internet access can only access the internet using a proxy server, use the following options when downloading the HCL:

Option	Description
<code>--proxyServer, --ps</code>	Provide the proxy server FQDN and port. For example: <code>--proxyServer proxy.example.com:3128</code> .
<code>--proxyHttps</code>	Add this option if the proxy server uses HTTPS.
<code>--proxyUser</code>	For a proxy server that requires authentication, enter the user name.
<code>--proxyPasswordFile</code>	For a proxy server that requires authentication, enter the path to a file where the password for proxy authentication is stored. The file content is used as the proxy password. For example, <code>--proxyPasswordFile ../../password.txt</code> .

Example that combines the options:

```
./lcm-bundle-transfer-util --vsanHclDownload --outputDirectory output-directory --  
proxyServer proxy.example.com:3128 --proxyUser vmwuser --proxyPasswordFile ../../  
password.txt --proxyHttps
```

1. Download the most recent version of the Bundle Transfer Utility on a computer with internet access.
 - a) Log in to the Broadcom Support Portal and browse to **My Downloads > VMware Cloud Foundation**.
 - b) Click the version of VMware Cloud Foundation to which you are upgrading.
 - c) Click **Drivers & Tools**.
 - d) Click the download icon for the Bundle Transfer Utility.
2. Extract `lcm-tools-prod.tar.gz`.

3. Navigate to the `lcm-tools-prod/bin/` and confirm that you have execute permission on all folders.
4. Copy the bundle transfer utility to a computer with access to the SDDC Manager appliance and then copy the bundle transfer utility to the SDDC Manager appliance.

- a) SSH in to the SDDC Manager appliance using the `vcf` user account.
- b) Enter `su` to switch to the root user.
- c) Create the `lcm-tools` directory.

```
mkdir /opt/vmware/vcf/lcm/lcm-tools
```

NOTE

If the `/opt/vmware/vcf/lcm/lcm-tools` directory already exists with an older version of the Bundle Transfer Utility, you need to delete contents of the existing directory before proceeding.

- d) Copy the Bundle Transfer Utility file (`lcm-tools-prod.tar.gz`) that you downloaded in step 1 to the `opt/vmware/vcf/lcm/lcm-tools` directory.
- e) Extract the contents of `lcm-tools-prod.tar.gz`.

```
tar -xvf lcm-tools-prod.tar.gz
```

- f) Set the permissions for the `lcm-tools` directory.

```
cd /opt/vmware/vcf/lcm/
chown vcf_lcm:vcf -R lcm-tools
chmod 750 -R lcm-tools
```

5. On the computer with internet access, download the HCL file.

```
./lcm-bundle-transfer-util --vsanHclDownload --outputDirectory output-directory
```

It can also be downloaded to the default path:

```
./lcm-bundle-transfer-util --vsanHclDownload
```

6. Copy the HCL file to the SDDC Manager appliance.
7. From the SDDC Manager appliance, use the Bundle Transfer Utility to upload the HCL file.

```
./lcm-bundle-transfer-util --vsanHclUpload --inputDirectory hcl-file-path --
sddcMgrFqdn sddc-manager-fqdn --sddcMgrUser user
```

<code>hcl-file-path</code>	Path from where HCL file should be picked up to upload. e.g <code>/root/testdownload/vsan/hcl/all.json</code> . If not given default will be taken. (<code>/root/PROD2/vsan/hcl/all.json</code>)
<code>sddc-manager-fqdn</code>	SDDC Manager FQDN. If not given default will be taken.
<code>user</code>	SDDC Manager user. After this, the tool will prompt for the user password.

Download Bundles to an Offline Depot

VMware Cloud Foundation 5.2 and later support an offline depot that you can connect to from multiple instances of SDDC Manager. Use the Bundle Transfer Utility to download and transfer bundles to the offline depot and then any SDDC Manager connected to the offline depot can access the bundles.

- [Set up an offline depot.](#)
- The offline depot must have:
 - The latest version of the Bundle Transfer Utility. You can download it from the Broadcom Support portal.
 - Internet connectivity (either directly or through a proxy) for downloading the bundles and other required files.
 - Configure TCP keepalive in your SSH client to prevent socket connection timeouts when using the Bundle Transfer Utility for long-running operations.
- Connect SDDC Manager to the offline depot. See [Connect SDDC Manager to a Software Depot for Downloading Bundles](#).

NOTE

You can also connect SDDC Manager to the offline depot after you download bundles to the offline depot.

You can use the Bundle Transfer Utility to download upgrade bundles and async patch bundles to the offline depot.

1. On the computer hosting offline depot, run the following command to download the bundles required to upgrade VMware Cloud Foundation.

```
./lcm-bundle-transfer-util --setUpOfflineDepot downloadPartnerBundle -sv vcf-source-version --offlineDepotRootDir offline-depot-root-dir --offlineDepotUrl url:port --depotUser user-name --depotUserPasswordFile path-to-password-file --partnerDepotUser user-name --partnerDepotUserPasswordFile path-to-password-file
```

For example:

```
./lcm-bundle-transfer-util --setUpOfflineDepot downloadPartnerBundle -sv 5.0.0.0 --offlineDepotRootDir /var/www --offlineDepotUrl https://10.123.456.78:8282 --depotUser user@example.com --depotUserPasswordFile ../vmw-depot --partnerDepotUser partner@example.com --partnerDepotUserPasswordFile ../partner-depot
```

2. Run the following command to download async patch bundles to the offline depot:

```
./lcm-bundle-transfer-util --setUpOfflineDepot --asyncPatches -offlineDepotRootDir offline-depot-root-dir --offlineDepotUrl url:port --depotUser user-name --depotUserPasswordFile path-to-password-file
```

For example:

```
./lcm-bundle-transfer-util --setUpOfflineDepot --asyncPatches -offlineDepotRootDir /var/www --offlineDepotUrl https://10.123.456.78:8282 --depotUser user@example.com --depotUserPasswordFile ../vmw-depot
```

After the bundles are available in the offline depot, you can use the SDDC Manager UI to apply the bundles to workload domains. Multiple instances of SDDC Manager UI can connect to the same offline depot.

VMware Cloud Foundation Upgrade Prerequisites

Before you upgrade VMware Cloud Foundation, make sure that the following prerequisites are met.

Table 241: Upgrade Prerequisites

Prerequisite	Additional Information
Allocate a temporary IP address for each vCenter Server upgrade	[Conditional] When upgrading from VMware Cloud Foundation 4.5.x.

Table continued on next page

Continued from previous page

Prerequisite	Additional Information
	Required for each vCenter Server upgrade. Must be allocated from the management subnet. The IP address can be reused.
Obtain updated licenses	New licenses required for: <ul style="list-style-type: none"> vSAN 8.x vSphere 8.x
Verify there are no expired or expiring passwords	Review the password management dashboard in SDDC Manager.
Verify there are no expired or expiring certificates	Review the Certificates tab in SDDC Manager for each workload domain.
Verify ESXi host TPM module status	[Conditional] If ESXi hosts have TPM modules in use, verify they are running the latest 2.0 firmware. If not in use they must be disabled in the BIOS. See KB 312159
Verify ESXi hardware is compatible with target version	See ESXi Requirements and VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php .
Manually update the vSAN HCL database to ensure that it is up-to-date.	See KB 2145116
Back up SDDC Manager, all vCenter Server instances, and NSX Manager instances.	Take file-based backups or image-level backups of SDDC Manager, all vCenter Servers, and NSX Managers. Take a cold snapshot of SDDC Manager.
Make sure that there are no failed workflows in your system and none of the VMware Cloud Foundation resources are in activating or error state.	CAUTION If any of these conditions are true, contact VMware Technical Support before starting the upgrade.
Review the <i>Release Notes</i> for known issues related to upgrades.	
Deactivate all VMware Cloud Foundation 4.x async patches and run an inventory sync before upgrading.	VMware Cloud Foundation 5.0 and later no longer require using the Async Patch Tool to enable upgrades from an async-patched VMware Cloud Foundation instance. See VCF Async Patch Tool Options for more information
Review Operational Impacts of NSX Upgrade in <i>NSX Upgrade Guide</i> to understand the impact that each component upgrade might have on your environment.	
In the vSphere Client, ensure there are no active alarms on hosts or vSphere clusters.	
Download the upgrade bundles.	See Downloading VMware Cloud Foundation Upgrade Bundles .

VMware Cloud Foundation 5.2.x Upgrade Overview

This section describes the tasks required to perform an upgrade to VMware Cloud Foundation 5.2.x.

VMware Cloud Foundation Upgrade Preparation

Review the [VMware Cloud Foundation Upgrade Prerequisites](#) before starting an upgrade.

Management Domain Upgrade

Table 242: SDDC Manager Upgrade

Task	Applies When	Additional Information
<ul style="list-style-type: none"> Precheck Update - Versions Prior to SDDC Manager 5.0 Perform Update Precheck in SDDC Manager 		
Apply the VMware Cloud Foundation Upgrade Bundle	<ul style="list-style-type: none"> The initial VMware Cloud Foundation version is <ul style="list-style-type: none"> – 4.5.x or 5.x 	If the current version of VMware Cloud Foundation is 4.5.x or 5.x Upgrade SDDC Manager to 5.2.x.
Apply the VMware Cloud Foundation Configuration Updates	<ul style="list-style-type: none"> Once the SDDC Manager has been upgraded to 5.2.x the Configuration updates can be applied collectively. 	
Update Compatibility Data with the Bundle Transfer Utility		[Conditional] Required when using offline bundle download

Table 243: Upgrade VMware Aria Suite

Task	Additional Information
Upgrade VMware Aria Suite Lifecycle for VMware Cloud Foundation	[Conditional] If VMware Aria Suite Lifecycle is present
Upgrade VMware Aria Suite products for VMware Cloud Foundation	[Conditional] If VMware Aria Suite products are present

Table 244: Upgrade NSX With Federation

Task	Applies When	Additional Information
Upgrade NSX Global Managers to 4.2	When NSX is deployed in the workload domain with NSX Federation configured.	<ul style="list-style-type: none"> [Conditional] If NSX Federation is present Upgrade NSX Global Managers to 4.2 using the Global Manager UI Upgrade standby global manager, followed by active global manager [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade all NSX global managers in your estate now.
Upgrade to NSX 4.2		<ul style="list-style-type: none"> Upgrade NSX to 4.2 using SDDC Manager [Optional] If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now.

Table continued on next page

Continued from previous page

Task	Applies When	Additional Information
		<ul style="list-style-type: none"> NSX upgrades across VI workload domains can be completed in sequence or up to five in parallel.

Table 245: Upgrade NSX Without Federation

Task	Applies When	Additional Information
Upgrade to NSX 4.2	When NSX is deployed in the workload domain and is not using NSX Federation.	<ul style="list-style-type: none"> Upgrade NSX to 4.2 using SDDC Manager. [Conditional] for VI Workload Domain upgrades. If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now.

Table 246: Upgrade vCenter Server

Task	Additional Information
Upgrade vCenter Server for VMware Cloud Foundation	<ul style="list-style-type: none"> [Conditional] When upgrading from VMware Cloud Foundation 4.5.x. Requires a temporary IP address in the management subnet [Conditional] When upgrading to VMware Cloud Foundation 5.2.1 using vCenter Reduced Downtime Upgrade (RDU). Requires a temporary IP address in the management subnet [Conditional] for VI Workload Domain upgrades. If you are upgrading by component rather than by workload domain, upgrade vCenter Servers that share a SSO Domain across all VI workload domains now in a serial order. Isolated Workload Domains can be upgraded in parallel

Table 247: Upgrade VxRail Manager and Management Domain vSphere clusters

Task	Additional Information
Upgrade vSAN Witness Host for VMware Cloud Foundation	[Conditional] If the vSphere cluster is a stretched vSAN cluster
Upgrade VxRail Manager and ESXi Hosts	<ul style="list-style-type: none"> Choose an approach based on your requirements. [Optional] If you are upgrading by component rather than by workload domain, upgrade vSphere clusters across all VI workload domains now.

Table 248: Post Upgrade Tasks

Task	Additional Information
Update Licenses for a Workload Domain	<p>[Conditional] If upgrading from a VMware Cloud Foundation version prior to 5.0</p> <p>Update licenses for:</p> <ul style="list-style-type: none"> • vSAN 8.x • vSphere 8.x
Apply Configuration Updates	[Conditional] If there are configuration updates required
Upgrade vSphere Distributed Switch versions	<ul style="list-style-type: none"> • [Optional] The upgrade lets the distributed switch take advantage of features that are available only in the later versions.
Upgrade vSAN on-disk format versions	<ul style="list-style-type: none"> • The upgrade lets the vSAN Cluster take advantage of features that are available only in the later versions. • The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure. • These updates can be performed at a time that is most convenient for your organization..

VI Workload Domain Upgrade**Table 249: Upgrade Precheck**

Task	Additional Information
Perform an upgrade precheck	

Table 250: Upgrade NSX Without Federation

Task	Applies When	Additional Information
Upgrade to NSX 4.2	When NSX is deployed in the workload domain and is not using NSX Federation.	<ul style="list-style-type: none"> • Upgrade NSX to 4.2 using SDDC Manager. • [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now.

Table 251: Upgrade NSX With Federation

Task	Applies When	Additional Information
Upgrade NSX Global Managers to 4.2	When NSX is deployed in the workload domain with NSX Federation configured.	<ul style="list-style-type: none"> • [Conditional] If NSX Federation is present

Table continued on next page

Continued from previous page

Task	Applies When	Additional Information
		<ul style="list-style-type: none"> Upgrade NSX Global Managers to 4.2 using the Global Manager UI Upgrade standby global manager, followed by active global manager [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade all NSX global managers in your estate now.
Upgrade to NSX 4.2		<ul style="list-style-type: none"> Upgrade NSX to 4.2 using SDDC Manager [Optional] If you are upgrading by component rather than by workload domain, upgrade NSX across all VI workload domains now. NSX upgrades across VI workload domains can be completed in sequence or up to five in parallel.

Table 252: Upgrade vCenter Server

Task	Additional Information
Upgrade vCenter Server for VMware Cloud Foundation	<ul style="list-style-type: none"> [Conditional] When upgrading from VMware Cloud Foundation 4.5.x. Requires a temporary IP address in the management subnet [Conditional] When upgrading to VMware Cloud Foundation 5.2.1 using vCenter Reduced Downtime Upgrade (RDU). Requires a temporary IP address in the management subnet [Conditional] for VI Workload Domain upgrades, If you are upgrading by component rather than by workload domain, upgrade vCenter Servers that share a SSO Domain across all VI workload domains now in a serial order. Isolated Workload Domains can be upgraded in parallel

Table 253: Upgrade VxRail Manager and VI Workload Domain vSphere clusters

Task	Additional Information
Upgrade vSAN Witness Host for VMware Cloud Foundation	[Conditional] If the vSphere cluster is a stretched vSAN cluster
Upgrade VxRail Manager and ESXi Hosts	<ul style="list-style-type: none"> Choose an approach based on your requirements.

Table continued on next page

Continued from previous page

Task	Additional Information
	<ul style="list-style-type: none"> • [Optional] If you are upgrading by component rather than by workload domain, upgrade vSphere clusters across all VI workload domains now.

Table 254: Post Upgrade Tasks

Task	Additional Information
Update Licenses for a Workload Domain	<p>[Conditional] If upgrading from a VMware Cloud Foundation version prior to 5.0</p> <p>Update licenses for:</p> <ul style="list-style-type: none"> • vSAN 8.x • vSphere 8.x
Apply Configuration Updates	[Conditional] If there are configuration updates required
Upgrade vSphere Distributed Switch versions	<ul style="list-style-type: none"> • [Optional] The upgrade lets the distributed switch take advantage of features that are available only in the later versions.
Upgrade vSAN on-disk format versions	<ul style="list-style-type: none"> • The upgrade lets the vSAN Cluster take advantage of features that are available only in the later versions. • The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure. • These updates can be performed at a time that is most convenient for your organization..

Upgrade the Management Domain to VMware Cloud Foundation 5.2.x

To upgrade to VMware Cloud Foundation 5.2.x, the management domain must be at VMware Cloud Foundation 4.5 or higher. If your environment is at a version lower than 4.5, you must upgrade the management domain to 4.5 or later and then upgrade to 5.2.x.

Until SDDC Manager is upgraded to version 5.2.x, you must upgrade the management domain before you upgrade VI workload domains. Once SDDC Manager is at version 5.2 or later, you can upgrade VI workload domains before or after upgrading the management domain, as long as all components in the workload domain are compatible.

Upgrade the components in the management domain in the following order:

1. SDDC Manager and VMware Cloud Foundation services.
2. VMware Aria Suite Lifecycle
3. NSX Manager and NSX Global Managers (if applicable).
4. vCenter Server.
5. VxRail Manager and ESXi.

After all upgrades have completed successfully:

1. Remove the VM snapshots you took before starting the update.
2. Take a backup of the newly installed components.

Perform Update Precheck - Versions Prior to SDDC Manager 5.0

If you have not yet upgraded to SDDC Manager 5.0, these are the steps to run a Precheck. You must perform a precheck before applying an update or upgrade bundle to ensure that your environment is ready for the update.

If you silence a vSAN Skyline Health alert in the vSphere Client, SDDC Manager skips the related precheck and indicates which precheck it skipped. Click **Restore Precheck** to include the silenced precheck. For example:

Hardware compatibility - SCSI controller is VMware certified

Description	Hardware compatibility - SCSI controller is VMware certified
Start Time	Sep 16, 2022, 10:44:11 AM
End Time	Sep 16, 2022, 10:44:12 AM
Health Status	Silenced RESTORE PRECHECK

You can also silence failed vSAN prechecks in the SDDC Manager UI by clicking **Silence Precheck**. Silenced prechecks do not trigger warnings or block upgrades.

IMPORTANT

You should only silence alerts if you know that they are incorrect. Do not silence alerts for real issues that require remediation.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the workload domain where you want to run the precheck.
3. On the domain summary page, click the **Updates/Patches** tab. The image below is a sample screenshot and may not reflect the correct product versions.

The screenshot shows the SDDC Manager interface for a workload domain. The left navigation pane includes Dashboard, Inventory, Workload Domains, Hosts, Repository, and Administration. The main content area shows the 'MGMT' management page with an 'ACTIVE' status. Resource usage is displayed as follows:

Resource	Total	Used	Free
CPU	76.61 GHz	16.19 GHz	60.42 GHz
Memory	468.75 GB	135.85 GB	332.9 GB
vSAN Storage	4.47 TB	0.6 TB	3.87 TB

The 'Update/Patches' tab is selected, showing a 'Precheck' section with a message: 'There is no recent precheck status available. It is recommended that you precheck your domain prior to update to prevent an error in an update.' A blue 'PRECHECK' button is located at the bottom right of this section.

4. Click **Precheck** to validate that the environment is ready to be upgraded.

Once the precheck begins, a message appears indicating the time at which the precheck was started.

The screenshot shows a 'Precheck' section with a dropdown arrow. Below it is a light blue message box containing an information icon, the text 'Precheck started at 1:40:08 PM, 06/19/2019 and is in progress.', and a 'View Status' link. Below the message box, there is a text prompt: 'You're currently prechecking your workload domain. It will take a while before it's done.' To the right of this text is a grey 'PRECHECK' button.

5. Click **View Status** to see detailed tasks and their status. The image below is a sample screenshot and may not reflect the correct versions.

The screenshot shows the 'Upgrade Precheck' section. At the top, a green message box states: 'Precheck passed at Jan 18, 2022, 2:52:30 PM. Resource health status will get updated in the next audit and applicable bundle will become available for upgrade.' Below this, a text prompt reads: 'It is recommended that you precheck your domain prior to update to prevent an error in an update'. To the right of this text are two buttons: a blue 'RUN PRECHECK' button and a grey 'PRECHECK FAILED RESOURCES' button. Below the text, there is a list of resources with expandable arrows and status indicators (green checkmarks or icons):

- ▼ SDDC Manager
 - > ✓ Lcm
 - > ✓ Operations Manager
 - > ✓ Multi Site Service
 - > ✓ Sddc Manager Ui
 - > ✓ Domain Manager
 - > ✓ Common Services
 - > ✓ VSAN - SDDC-Cluster1
 - > ✓ vip-nsx-mgmt.vrack.vsphere.local
 - > ✓ vcenter-1.vrack.vsphere.local
 - > ✓ PSC
- ▼ 4 ESXi Host
 - ▼ SDDC-Cluster1
 - > ✓ esxi-4
 - > ✓ esxi-3
 - > ✓ esxi-2
 - > ✓ esxi-1

6. To see details for a task, click the Expand arrow.

If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **Precheck Failed Resources** to retry all failed tasks.

7. If the workload domain contains a host that includes pinned VMs, the precheck fails at the Enter Maintenance Mode step. If the host can enter maintenance mode through vCenter Server UI, you can suppress this check for NSX and ESXi in VMware Cloud Foundation by following the steps below.
 1. Log in to SDDC Manager by using a Secure Shell (SSH) client with the user name vcf and password you specified in the deployment parameter workbook.
 2. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
 3. Add the following line to the end of the file:
`lcm.nsx.suppress.dry.run.emm.check=true`

`lcm.esx.suppress.dry.run.emm.check.failures=true`
 4. Restart Lifecycle Management by typing the following command in the console window.
`systemctl restart lcm`
 5. After Lifecycle Management is restarted, run the precheck again.

The precheck result is displayed at the top of the Upgrade Precheck Details window. If you click **Exit Details**, the precheck result is displayed at the top of the Precheck section in the Updates/Patches tab.

Ensure that the precheck results are green before proceeding. A failed precheck may cause the update to fail.


Perform Update Precheck in SDDC Manager

You must perform a precheck in SDDC Manager before applying an update bundle to ensure that your environment is ready for the update.

Bundle-level pre-checks for vCenter are available in VMware Cloud Foundation.

If you silence a vSAN Skyline Health alert in the vSphere Client, SDDC Manager skips the related precheck and indicates which precheck it skipped. Click **RESTORE PRECHECK** to include the silenced precheck. For example:

RESTORE PRECHECK

Resource Name	Description	Health Status	Error Description	Impact	Remediation
SDDC-Cluster1	Checks the age of the VMware Hardware Compatibility Guide database used for the HCL checks. Shows warning or error when it is older than 90 or 180 days, respectively. VMware updates the VCG frequently, so it is important to keep the local copy up-to-date	 SILENCED	Check skipped because 'com.vmware.vsan.health.test.hcldbuptodate' is silenced in vSAN		If you would like to run the check, click the Restore button or enable it through the vSphere UI

You can also silence failed vSAN prechecks in the SDDC Manager UI by clicking **Silence Precheck**. Silenced prechecks do not trigger warnings or block upgrades.

IMPORTANT

Only silence alerts if you know that they are incorrect. Do not silence alerts for real issues that require remediation.

1. In the navigation pane, click **Inventory** > **Workload Domains**.

2. On the Workload Domains page, click the workload domain where you want to run the precheck.
3. On the domain summary page, click the **Updates** tab.

(The following image is a sample screenshot and may not reflect current product versions.)

The screenshot shows the 'Precheck Status' section for workload domain 'sddcId-1001'. The status is 'Precheck completed successfully' and was completed on Dec 21, 2022, at 10:04:47 AM. There are 0 silenced prechecks for the entire domain. A 'RUN PRECHECK' button is located in the top right corner of the status box.

NOTE

It is recommended that you Precheck your workload domain prior to performing an upgrade.

4. Click **RUN PRECHECK** to select the components in the workload domain you want to precheck.
 - a) You can select to run a Precheck only on vCenter or the vSphere cluster. All components in the workload domain are selected by default. To perform a precheck on certain components, choose **Custom selection**.

The screenshot shows the 'Precheck' configuration dialog. The 'Run Precheck on' section is set to 'Entire Workload Domain'. The 'Target Version' is 'General Upgrade Readiness'. A table lists components to be prechecked:

Component	Description
<input checked="" type="checkbox"/> sddc-manager.vrack.vsphere.local	SDDC_MANAGER
<input checked="" type="checkbox"/> vip-nsx-mgmt.vrack.vsphere.local	NSX
<input checked="" type="checkbox"/> SDDC-Cluster1	CLUSTER
<input checked="" type="checkbox"/> vcenter-1.vrack.vsphere.local	VCENTER
<input checked="" type="checkbox"/> 4	4 objects

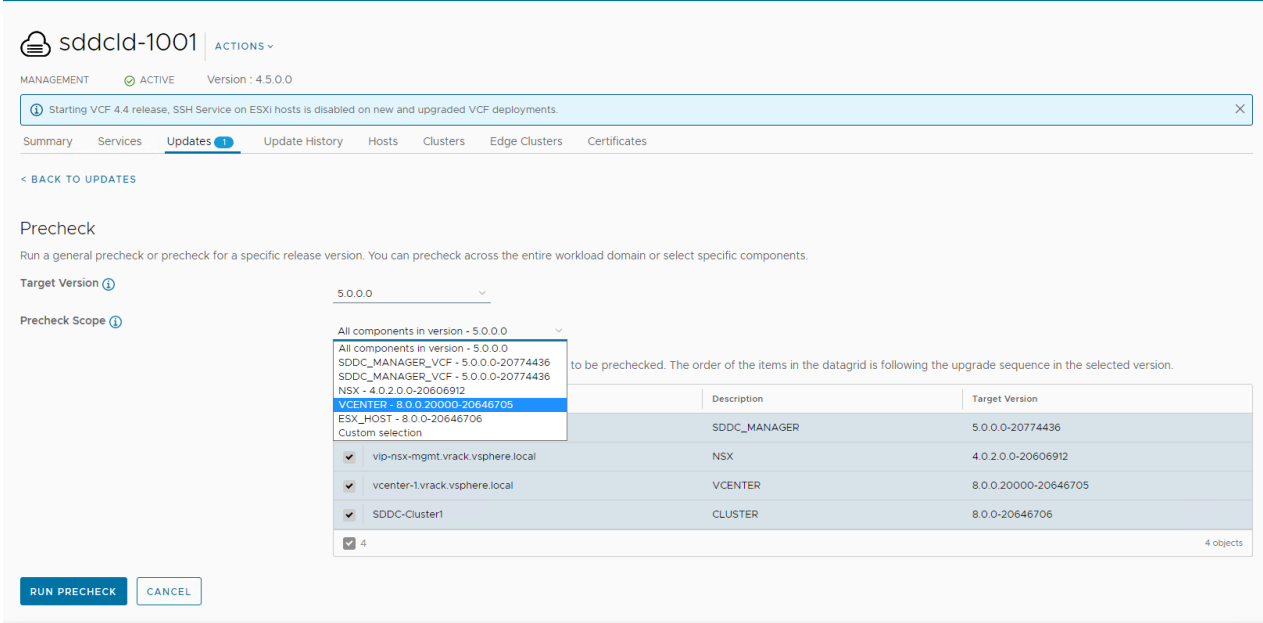
Buttons for 'RUN PRECHECK' and 'CANCEL' are visible at the bottom.

NOTE

For VMware Cloud Foundation on Dell EMC VxRail, you can run prechecks on VxRail Manager.

- b) If there are pending upgrade bundles available, then the "Target Version" dropdown contains "General Upgrade Readiness" and the available VMware Cloud Foundation versions to upgrade to. If there is an

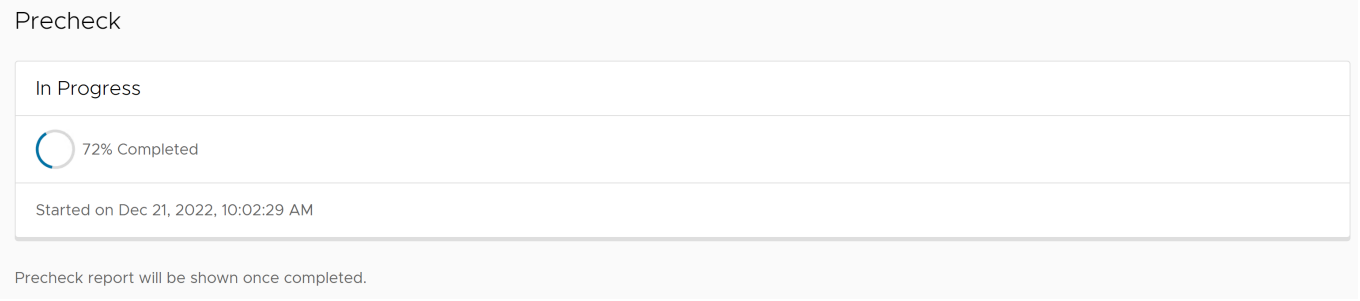
available VMware Cloud Foundation upgrade version, there will be extra checks - bundle-level prechecks for hosts, vCenter Server, and so forth. The version specific prechecks will only run prechecks on components that have available upgrade bundles downloaded.



The screenshot shows the 'Precheck' configuration interface in the VMware Cloud Foundation management console. The target version is set to 5.0.0.0. The precheck scope is set to 'All components in version - 5.0.0.0'. A table lists the components to be prechecked, including SDDC_MANAGER, NSX, vcenter-1, and SDDC-Cluster1. The table has columns for Description and Target Version.

Description	Target Version
SDDC_MANAGER	5.0.0.0-20774436
NSX	4.0.2.0.0-20606912
vcenter-1.vrack.vsphere.local	8.0.0.20000-20646705
SDDC-Cluster1	8.0.0-20646706

- When the precheck begins, a progress message appears indicating the precheck progress and the time when the precheck began.



The screenshot shows the 'Precheck' progress screen in the VMware Cloud Foundation management console. The progress is 72% completed, and the precheck started on Dec 21, 2022, at 10:02:29 AM. The screen displays 'In Progress' and a progress indicator.

NOTE

Parallel precheck workflows are supported. If you want to precheck multiple domains, you can repeat steps 1-5 for each of them without waiting for step 5 to finish.

- Once the Precheck is complete, the report appears. Click through **ALL**, **ERRORS**, **WARNINGS**, and **SILENCED** to filter and browse through the results.

Precheck

Results Completed on Dec 21, 2022, 10:04:47 AM

✔ 174 Passed
❌ 7 Errors
⚠ 3 Warnings
🔇 0 Silenced


[RETRY ALL FAILED RESOURCES](#)

Report

[ALL](#)
[ERRORS](#)
[WARNINGS](#)
[SILENCED](#)

Version: General Upgrade Readiness (4 component)

- > [SDDC Manager sddc-manager.vrack.vsphere.local](#)
- > [NSX vip-nsx-mgmt.vrack.vsphere.local](#)
- > [ESXi Host Cluster SDDC-Cluster1](#)
- > [vCenter vcenter-1.vrack.vsphere.local](#)

Resource Name	Description	Health Status	Error Description	Impact	Remediation
 Precheck entry not selected					

7. To see details for a task, click the expander arrow.

If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **RETRY ALL FAILED RESOURCES** to retry all failed tasks.

8. If the workload domain contains a host that includes pinned VMs, the precheck fails at the Enter Maintenance Mode step. If the host can enter maintenance mode through vCenter Server UI, you can suppress this check for NSX and ESXi in VMware Cloud Foundation by following the steps below.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client with the user name `vcf` and password.
2. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
3. Add the following line to the end of the file:

```
lcm.nsxt.suppress.dry.run.emm.check=true
```

```
lcm.esx.suppress.dry.run.emm.check.failures=true
```

4. Restart Lifecycle Management by typing the following command in the console window.

```
systemctl restart lcm
```

5. After Lifecycle Management is restarted, run the precheck again.

The precheck result is displayed at the top of the Upgrade Precheck Details window. If you click **Exit Details**, the precheck result is displayed at the top of the Precheck section in the Updates tab.

Ensure that the precheck results are green before proceeding. Although a failed precheck will not prevent the upgrade from proceeding, it may cause the update to fail.

Apply the VMware Cloud Foundation 5.2.x Upgrade Bundle

The VMware Cloud Foundation Upgrade bundle upgrades the SDDC Manager appliance and VMware Cloud Foundation services.

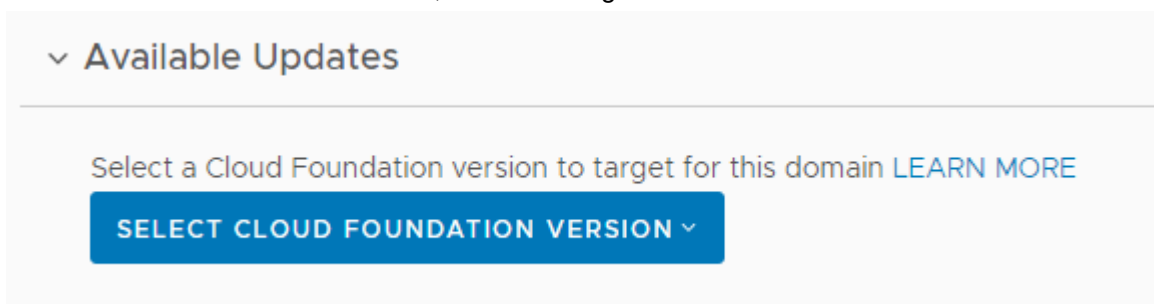
- Download the VMware Cloud Foundation update bundle for your target release. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Ensure you have a recent successful backup of SDDC Manager using an external SFTP server.
- Ensure you have taken a snapshot of the SDDC Manager appliance.
- Ensure you have recent successful backups of the components managed by SDDC Manager.
- [Perform Update Precheck in SDDC Manager](#) and resolve any issues.

After SDDC Manager is upgraded to 5.2 or later, new functionality is introduced that allows you to upgrade SDDC Manager without having to upgrade the entire VMware Cloud Foundation BOM. See [Independent SDDC Manager Upgrade using the SDDC Manager UI](#).

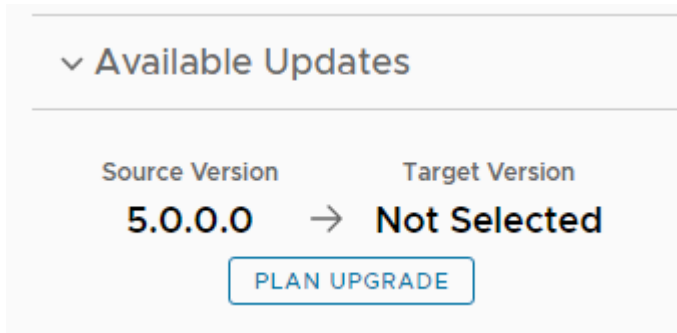
1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. On the Workload Domains page, click the management domain and then click the **Updates** tab.
3. In the **Available Updates** section, select the target VMware Cloud Foundation release or click **Plan Upgrade**.

The available options depend on the source version of VMware Cloud Foundation.

- For VMware Cloud Foundation 4.5.x, select the target version.



- For VMware Cloud Foundation 5.x, click **Plan Upgrade**, select a target version, and click **Confirm**.



4. Click **Update Now** or **Schedule Update** next to the VMware Cloud Foundation Upgrade bundle.
5. If you selected **Schedule Update**, select the date and time for the bundle to be applied and click **Schedule**.

Schedule Update ✕

The bundle will be scheduled based on your selected date and time.

Date 

Time

CANCEL

SCHEDULE

If you clicked **Update Now**, the VMware Cloud Foundation Update Status window displays the components that will be upgraded and the upgrade status. Click **View Update Activity** to view the detailed tasks. After the upgrade is completed, a green bar with a check mark is displayed.

6. Click **Finish**.

When the update completes successfully, you are logged out of the SDDC Manager UI and must log in again.

Apply VMware Cloud Foundation Configuration Updates

VMware Cloud Foundation Configuration Updates identifies and resolves any discrepancies between the intended/prescribed configuration and the actual configuration, ensuring that the deployment aligns with the recommended configuration. This process includes reconciling the configuration for 2nd party software components listed in the VMware Cloud Foundation Bill of Materials (BOM).

Configuration updates may be required after you apply software updates. Once a configuration update becomes available, you can apply it immediately or wait until after you have applied all software updates. Configuration Updates must be performed during a maintenance window.

Configuration Updates can be applied to multiple domains in parallel. However, if a Configuration Update is in progress, another configuration update on the same domain should not be attempted.

NOTE

Configuration Updates in VCF detects and reconciles to a prescribed configuration for the release. Once reconciled, it does not identify subsequent non-compliance arising from out of band changes.

The following configuration updates may become available, depending on your source version of VMware Cloud Foundation:

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
ConfigureVsanHalsolutionAddressesConfigDrift	Configures the vSAN HA network isolation address to use the vSAN vmkernel interface gateway, in conformance with VCF best practices.	4.3.0.0	CLUSTER	FIX	vCenter 7.0.3

Table continued on next page

Continued from previous page

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
ToggleVSanRecommendationConfigDrift	Disables vSAN baseline recommendations for vSAN enabled clusters.	4.4.1.0	CLUSTER	FIX	vCenter 7.0.0
RemoveNfsDatastoreConfigDrift	Removes NFS datastore on hosts.	5.0.0.0	CLUSTER	FIX	NA
CloudAdminRoleConfigDrift	Creates Cloud Admin role in vCenter Server for the management domain.	5.0.0.0	DOMAIN	FEATURE	vCenter 7.0.3
AllowBrokerConfigurationConfigDrift	Adds <code>config.SDDC.Deployed.AllowBrokerConfiguration</code> advanced property in vCenter Server. This property restricts the user from configuring an external IDP from the vCenter UI in the ELM ring (workload domain vCenters). Configuration is only possible from the management domain vCenter UI and isolated workload domain vCenter UI.	5.1.0.0	DOMAIN	FEATURE	vCenter 8.0.2
ClusterHaSettingsConfigDrift	Removes <code>das.includeFTComplianceChecks</code> option HA configuration from all clusters on the management domain.	5.1.0.0	DOMAIN	FEATURE	vCenter 8.0.1
ComputeManagerSettingsDrift	Creates an internal NSX service account to enable NSX to	5.1.0.0	DOMAIN	FEATURE	vCenter 7.0.2.00400, NSX 3.1.3.0.0

Table continued on next page

Continued from previous page

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
	vSphere Lifecycle Manager communication.				
DvpgConfigurationDrift	<p>Creates a new distributed virtual port group named VM_MANAGEMENT in the target domain, and migrates all VMs connected to the management port group to this new port group. The purpose of this feature is to allow separation of traffic coming from management VMs and ESXi hosts.</p> <p>VMs migrated: VCSA, SDDC Manager, NSX Manager and Edge VMs.</p>	5.1.0.0	CLUSTER	FEATURE	NA
EsxAdvancedOptionsConfigDrift	<p>Configures <code>UserVars.SuppressShellWarning</code> property on every ESXi host to false, to enable warnings for ESXi Shell and SSH services.</p>	5.1.0.0	DOMAIN	FEATURE	NA
WorkspaceOneBrokerConfigDrift	Configures BOM components as OIDC relying parties of Workspace ONE Broker in vCenter.	5.1.0.0	DOMAIN	FEATURE	vCenter 8.0.2, NSX 4.1.2
RegisterSDDCmanagerAsVCExtensionConfigDrift	Register SDDC Manager as an extension in a workload domain vCenter.	5.2.0.0	DOMAIN	FEATURE	vCenter 7.0.0

Table continued on next page

Continued from previous page

Configuration Update	Description	Introduced in VCF Version	Resource Type	Update Type	Required Minimum Component Versions
SddcMgrVxRailServiceAccountConfigDrift	Creates a service account for SDDC Manager to VxRail Manager communication.	5.2.0.0	CLUSTER	FEATURE	vCenter 7.0.400

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the workload domain name and then click the **Updates** tab.
3. Click **Run Precheck** to run the upgrade precheck.
Resolve any issues before proceeding with the upgrade.
4. Expand **Available Configuration Updates**, click **Apply All**.

Available Configuration Updates 4

Configuration updates may be required after you apply software updates. Once a configuration update becomes available, you can apply it immediately or wait until after you have applied all software updates.

[APPLY ALL](#)

Description	Type	Resource Type
Remove das.includeFTcomplianceChecks option from all clusters HA configuration	FEATURE	DOMAIN
ESXi advanced options for ESXi Hosts Addition Drift	FEATURE	DOMAIN
Register SDDC Manager as an extension in a domain vCenter	FEATURE	DOMAIN
Creates a Distributed Virtual Port Group to enable traffic isolation between management VMs and ESXi hosts	FEATURE	CLUSTER

- **FEATURE:** Configuration change required for a new feature.
- **FIX:** Configuration change associated with a fix for a defect.

5. Check the progress of a configuration update by clicking the task in the Tasks panel.

Tasks ⊞ ↻ ↗ ✕

[REFRESH](#) [RESET FILTERS](#)

Task	Subtask	Task Status	Last Occurrence
Configuration Updates	DVPG Configuration Drift	<div style="width: 50%;"><div style="background-color: #0070C0; height: 10px;"></div></div> 50%	5/19/23, 10:55 AM

6. After the configuration updates are successfully applied, they will no longer appear in the table.

Available Configuration Updates

✔ The recent configuration update successfully completed, please check upgrade details. [View Details](#)

No Available Updates.

Pending Configuration Updates do not block future BOM upgrades.

Upgrade VMware Aria Suite Lifecycle and VMware Aria Suite Products for VMware Cloud Foundation

VMware Cloud Foundation does not manage upgrades for VMware Aria Suite Lifecycle and the VMware Aria Suite products. Use VMware Aria Suite Lifecycle to upgrade VMware Aria Suite products.

If you had VMware Aria Suite Lifecycle, VMware Aria Operations for Logs, VMware Aria Automation, VMware Aria Operations, or Workspace ONE Access in your pre-upgrade environment, you must upgrade them from VMware Aria Suite Lifecycle.

Use VMware Aria Suite Lifecycle to:

- Download upgrade binaries
- Create snapshots of the virtual appliances
- Run pre-upgrade checks
- Upgrade VMware Aria Suite products

You can upgrade VMware Aria Suite products as new versions become available in VMware Aria Suite Lifecycle. VMware Aria Suite Lifecycle will only allow upgrades to compatible and supported versions of VMware Aria Suite products.

NOTE

See the [VMware Interoperability Matrix](#) for information about which versions are supported with your version of VMware Cloud Foundation and [KB 88829](#) for more information about supported upgrade paths using VMware Aria Suite Lifecycle.

IMPORTANT

The VMware Cloud Foundation 5.2 BOM requires VMware Aria Suite Lifecycle 8.18 or higher.

NOTE

The VMware Aria Suite of products were formerly known as the vRealize Suite of products.

1. Log in to VMware Aria Suite Lifecycle at `https://<aria_suite_lifecycle_manager_fqdn>` as the administrator.
2. Upgrade VMware Aria Suite products.

Upgrade VMware Aria Suite Lifecycle first and then upgrade VMware Aria Suite products.

See “Upgrading VMware Aria Suite Lifecycle and VMware Aria Suite Products” in the *VMware Aria Suite Lifecycle Installation, Upgrade, and Management Guide* for your current version of [VMware Aria Suite Lifecycle](#).

Upgrade NSX for VMware Cloud Foundation in a Federated Environment

If NSX Federation is configured between two VMware Cloud Foundation instances, SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must manually upgrade the NSX Global Managers for each instance.

Download NSX Global Manager Upgrade Bundle

SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must download the NSX upgrade bundle manually to upgrade the NSX Global Managers.

1. Log in to the Broadcom Support Portal and browse to **My Downloads > VMware NSX**.
2. Click the version of NSX to which you are upgrading.
3. Locate the **NSX version Upgrade Bundle** and verify that the upgrade bundle filename extension ends with `.mub`.

The upgrade bundle filename has the following format `VMware-NSX-upgrade-bundle-versionnumber.buildnumber.mub`.

4. Click the download icon to download the upgrade bundle to the system where you access the NSX Global Manager UI.

Upgrade the Upgrade Coordinator for NSX Federation

The upgrade coordinator runs in the NSX Manager. It is a self-contained web application that orchestrates the upgrade process of hosts, NSX Edge cluster, NSX Controller cluster, and the management plane.

The upgrade coordinator guides you through the upgrade sequence. You can track the upgrade process and, if necessary, you can pause and resume the upgrade process from the UI.

1. In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Proceed to Upgrade**.
4. Navigate to the upgrade bundle .mub file you downloaded or paste the download URL link.
 - Click **Browse** to navigate to the location you downloaded the upgrade bundle file.
 - Paste the VMware download portal URL where the upgrade bundle .mub file is located.
5. Click **Upload**.
When the file is uploaded, the **Begin Upgrade** button appears.
6. Click **Begin Upgrade** to upgrade the upgrade coordinator.

NOTE

Upgrade one upgrade coordinator at a time.

7. Read and accept the EULA terms and accept the notification to upgrade the upgrade coordinator..
8. Click **Run Pre-Checks** to verify that all NSX components are ready for upgrade.
The pre-check checks for component connectivity, version compatibility, and component status.
9. Resolve any warning notifications to avoid problems during the upgrade.

Upgrade NSX Global Managers for VMware Cloud Foundation

Manually upgrade the NSX Global Managers when NSX Federation is configured between two VMware Cloud Foundation instances.

Before you can upgrade NSX Global Managers, you must upgrade all VMware Cloud Foundation instances in the NSX Federation, including NSX Local Managers, using SDDC Manager.

1. In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Start** to upgrade the management plane and then click **Accept**.
4. On the Select Upgrade Plan page, select **Plan Your Upgrade** and click **Next**.

The NSX Manager UI, API, and CLI are not accessible until the upgrade finishes and the management plane is restarted.

Upgrade NSX for VMware Cloud Foundation 5.2.x

Upgrade NSX in the management domain and VI workload domains. VMware Cloud Foundation 5.2.1 supports in-place host upgrades for clusters that use vSphere Lifecycle Manager baselines.

Until SDDC Manager is upgraded to version 5.2, you must upgrade NSX in the management domain before you upgrade NSX in a VI workload domain. Once SDDC Manager is at version 5.2 or later, you can upgrade NSX in VI workload domains before or after upgrading NSX in the management domain.

Upgrading NSX involves the following components:

- Upgrade Coordinator
- NSX Edges/Clusters (if deployed)
- Host clusters
- NSX Manager cluster

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates/Patches** tab.

When you upgrade NSX components for a selected VI workload domain, those components are upgraded for all VI workload domains that share the NSX Manager cluster.

3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

NOTE

The NSX precheck runs on all VI workload domains in your environment that share the NSX Manager cluster.

4. For VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for NSX.

The screenshot shows the 'Available Updates' section in the VMware Cloud Foundation interface. It displays the target version as 5.2.0.0 and indicates that 1 of 4 steps is done. There are buttons for 'VIEW DETAILS', 'SCHEDULE UPDATE', and 'UPDATE NOW'. Below this, it shows the release date (Jun 19, 2024) and size (9 GB) of the VMware Software Update 5.2.0.0. A note mentions that the upgrade bundle is for VMware NSX Data Center 4.2.0.0 and encourages running the NSX Upgrade Evaluation Tool. A link to the release notes is provided.

- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.
By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.
- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.
By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) If you selected the **Schedule Upgrade** option, specify the date and time for the NSX bundle to be applied and click **Next**.
- f) On the Review page, review your settings and click **Finish**.

If you selected **Upgrade Now**, the NSX upgrade begins and the upgrade components are displayed. The upgrade view displayed here pertains to the workload domain where you applied the bundle. Click the link to the associated workload domains to see the components pertaining to those workload domains. If you selected **Schedule Upgrade**, the upgrade begins at the time and date you specified.

5. For VMware Cloud Foundation 5.2.1:

- a) In the Available Updates section, click the **Configure Update** button.

- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.

By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.

- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.

By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.

By default ESXi hosts are placed into maintenance mode during an upgrade. Starting with VMware Cloud Foundation 5.2.1, in-place upgrades are available for workload domains in which all the clusters use vSphere Lifecycle Manager baselines. If NSX Manager is shared between workload domains, in-place upgrade is only available if all the clusters in all the workload domains that share the NSX Manager use vLCM baselines. If the option is available, you can select **In-place** as the upgrade mode to avoid powering off and placing hosts into maintenance mode before the upgrade.

NOTE

To perform an in-place upgrade, the target NSX version must be the VMware Cloud Foundation 5.2.1 BOM version or later.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) On the Review page, review your settings and click **Run Precheck**.

The precheck begins. Resolve any issues until the precheck succeeds.

- f) After the precheck succeeds, click **Schedule Update** and select an option.
6. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).

If a component upgrade fails, the failure is displayed across all associated workload domains. Resolve the issue and retry the failed task.

When all NSX workload components are upgraded successfully, a message with a green background and check mark is displayed.

Upgrade vCenter Server for VMware Cloud Foundation 5.2.x

The upgrade bundle for VMware vCenter Server is used to upgrade the vCenter Server instances managed by SDDC Manager. Upgrade vCenter Server in the management domain before upgrading vCenter Server in VI workload domains.

- Download the VMware vCenter Server upgrade bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Take a file-based backup of the vCenter Server appliance before starting the upgrade. See [Manually Back Up vCenter Server](#).

NOTE

After taking a backup, do not make any changes to the vCenter Server inventory or settings until the upgrade completes successfully.

- If your workload domain contains Workload Management (vSphere with Tanzu) enabled clusters, the supported target release depends on the version of Kubernetes (K8s) currently running in the cluster. Older versions of K8s might require a specific upgrade sequence. See [KB 92227](#) for more information.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

4. Upgrading to VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for vCenter Server.
 - b) Click **Confirm** to confirm that you have taken a file-based backup of the vCenter Server appliance before starting the upgrade.
 - c) If you selected **Schedule Update**, click the date and time for the bundle to be applied and click **Schedule**.
 - d) If you are upgrading from VMware Cloud Foundation 4.5.x, enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.

Schedule Update

- 1 Introduction
- 2 Temporary Network
- 3 Review

Temporary Network ×

The migration-based upgrade requires a temporary address to be used for a parallel instance. This will only be used during the upgrade.

IP Address

Subnet Mask

Gateway

- e) Review the upgrade settings and click **Finish**.
5. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 5.x:
- a) In the Available Updates section, click **Configure Update**.
 - b) Select the upgrade mechanism and click **Next**.

Option	Description
vCenter Reduced Downtime Upgrade	<p>The reduced downtime upgrade process uses a migration-based approach. In this approach, a new vCenter Server Appliance is deployed and the current vCenter data and configuration is copied to it.</p> <p>During the preparation phase of a reduced downtime upgrade, the source vCenter Server Appliance and all resources remain online. The only downtime occurs when the source vCenter Server Appliance is stopped, the configuration is switched over to the target vCenter, and the services are started. The downtime is expected to take approximately 5 minutes under ideal network, CPU, memory, and storage provisioning.</p> <p style="text-align: center;">NOTE To perform a vCenter Reduced Downtime Upgrade, the target vCenter version must be the VMware Cloud Foundation 5.2.1 BOM version or later.</p>

Table continued on next page

Continued from previous page

Option	Description
vCenter Regular Upgrade	During a regular upgrade, the vCenter Server Appliance is offline for the duration of the upgrade.

- c) Select a backup option and click **Next**.
- d) For an RDU update, provide a temporary network to be used only during the upgrade and click **Next**.

Automatic	Automatically assign network information.
Static	Enter an IP address, subnet mask, and gateway. The IP address must be in the management subnet.

- e) Schedule the update and click **Next**.

For vCenter Reduced Downtime Upgrade	Select scheduling options for the preparation and switchover phases of the upgrade. NOTE If you are scheduling the switchover phase, you must allow a minimum of 4 hours between the start of preparation and the start of switchover.
For vCenter Regular Upgrade	Select an Upgrade Now or Schedule Update .

- f) Review the upgrade settings and click **Finish**.
6. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 4.5.x:
 - a) In the Available Updates section, click **Configure Update**.
 - b) Enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.
 - c) Select a backup option and click **Next**.
 - d) Schedule the update and click **Next**.
 - e) Review the upgrade settings and click **Finish**.
 7. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).
 8. After the upgrade is complete, remove the old vCenter Server appliance (if applicable).

NOTE

Removing the old vCenter is only required for major upgrades. If you performed a vCenter RDU patch upgrade, the old vCenter is automatically removed after a successful upgrade.

If the upgrade fails, resolve the issue and retry the failed task. If you cannot resolve the issue, restore vCenter Server using the file-based backup. See [Restore vCenter Server](#). vCenter RDU upgrades perform automatic rollback if the upgrade fails.

Once the upgrade successfully completes, use the vSphere Client to change the vSphere DRS Automation Level setting back to the original value (before you took a file-based backup) for each vSphere cluster that is managed by the vCenter Server. See [KB 87631](#) for information about using VMware PowerCLI to change the vSphere DRS Automation Level.

Upgrade VxRail Manager and ESXi Hosts for VMware Cloud Foundation

Use the VxRail upgrade bundle to upgrade VxRail Manager and the ESXi hosts in the workload domain. Upgrade the management domain first and then VI workload domains.

- Validate that the ESXi passwords are valid.
- Download the VxRail upgrade bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.

By default, the upgrade process upgrades the ESXi hosts in all clusters in a workload domain in parallel. If you have multiple clusters in the management domain or in a VI workload domain, you can select the clusters to upgrade. You can also choose to upgrade the clusters in parallel or sequentially.

If you are using external (non-vSAN) storage, the following procedure updates the ESXi hosts attached to the external storage. However, updating and patching the storage software and drivers is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability after an ESXi upgrade, consult the vSphere HCL and your storage vendor.

1. Navigate to the **Updates/Patches** tab of the workload domain.
2. Click **Precheck** to run the upgrade precheck.
Resolve any issues before proceeding with the upgrade.
3. In the Available Updates section, select the target release.
4. Click **Upgrade Now** or **Schedule Update**.
If you selected **Schedule Update**, specify the date and time for the bundle to be applied.
5. Select the clusters to upgrade and click **Next**.
The default setting is to upgrade all clusters. To upgrade specific clusters, click **Enable cluster-level selection** and select the clusters to upgrade.
6. Click **Next**.
7. Select the upgrade options and click **Finish**.
By default, the selected clusters are upgraded in parallel. If you selected more than five clusters to be upgraded, the first five are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Click **Enable Quick Boot** if desired. Quick Boot for ESXi hosts is an option that allows Update Manager to reduce the upgrade time by skipping the physical reboot of the host.
8. Monitor the upgrade progress. See [Monitor VMware Cloud Foundation Updates](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [Upgrade vSAN on-disk format versions](#).

Upgrade vSAN Witness Host for VMware Cloud Foundation

If your VMware Cloud Foundation environment contains stretched clusters, update and remediate the vSAN witness host.

Download the ESXi ISO that matches the version listed in the the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes*.

1. In a web browser, log in to vCenter Server at `https://vcenter_server_fqdn/ui`.
2. Upload the ESXi ISO image file to vSphere Lifecycle Manager.

- a) Click **Menu** › **Lifecycle Manager**.
 - b) Click the **Imported ISOs** tab.
 - c) Click **Import ISO** and then click **Browse**.
 - d) Navigate to the ESXi ISO file you downloaded and click **Open**.
 - e) After the file is imported, click **Close**.
3. Create a baseline for the ESXi image.
- a) On the Imported ISOs tab, select the ISO file that you imported, and click **New baseline**.
 - b) Enter a name for the baseline and specify the **Content Type** as Upgrade.
 - c) Click **Next**.
 - d) Select the ISO file you had imported and click **Next**.
 - e) Review the details and click **Finish**.
4. Attach the baseline to the vSAN witness host.
- a) Click **Menu** › **Hosts and Clusters**.
 - b) In the Inventory panel, click **vCenter** › **Datacenter**.
 - c) Select the vSAN witness host and click the **Updates** tab.
 - d) Under Attached Baselines, click **Attach** › **Attach Baseline or Baseline Group**.
 - e) Select the baseline that you had created in step 3 and click **Attach**.
 - f) Click **Check Compliance**.
- After the compliance check is completed, the **Status** column for the baseline is displayed as Non-Compliant.
5. Remediate the vSAN witness host and update the ESXi hosts that it contains.
- a) Right-click the vSAN witness and click **Maintenance Mode** › **Enter Maintenance Mode**.
 - b) Click **OK**.
 - c) Click the **Updates** tab.
 - d) Select the baseline that you had created in step 3 and click **Remediate**.
 - e) In the End user license agreement dialog box, select the check box and click **OK**.
 - f) In the Remediate dialog box, select the vSAN witness host, and click **Remediate**.
- The remediation process might take several minutes. After the remediation is completed, the **Status** column for the baseline is displayed as Compliant.
- g) Right-click the vSAN witness host and click **Maintenance Mode** › **Exit Maintenance Mode**.
 - h) Click **OK**.

Upgrade vSphere Distributed Switch versions

[Optional] Upgrade the distributed switch to take advantage of features that are available only in the later versions.

ESXi and vCenter Upgrades are completed.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Upgrade** › **Upgrade Distributed Switch**
3. Select the vSphere Distributed Switch version that you want to upgrade the switch to and click **Next**

The vSphere Distributed Switch is successfully upgraded.

Upgrade vSAN on-disk format versions

[Optional] Upgrade the vSAN on-disk format version to take advantage of features that are available only in the later versions.

- ESXi and vCenter Upgrades are completed
- Verify that the disks are in a healthy state. Navigate to the Disk Management page to verify the object status.
- Verify that your hosts are not in maintenance mode. When upgrading the disk format, do not place the hosts in maintenance mode.
- Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster. For information about vSAN resynchronization, see vSphere Monitoring and Performance
- The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure.

1. Navigate to the vSAN cluster.
2. Click the **Configure** tab.
3. Under **vSAN**, select **Disk Management**.
4. Click **Pre-check Upgrade**. The upgrade pre-check analyzes the cluster to uncover any issues that might prevent a successful upgrade. Some of the items checked are host status, disk status, network status, and object status. Upgrade issues are displayed in the **Disk pre-check status text** box.
5. Click **Upgrade**.
6. Click **Yes** on the Upgrade dialog box to perform the upgrade of the on-disk format.

vSAN successfully upgrades the on-disk format. The On-disk Format Version column displays the disk format version of storage devices in the cluster

Update License Keys for a Workload Domain

If upgrading from a VMware Cloud Foundation version prior to 5.0, you need to update your license keys to support vSAN 8.x and vSphere 8.x.

You need a new license key for vSAN 8.x and vSphere 8.x. Prior to VMware Cloud Foundation 5.1.1, you must add and update the component license key for each upgraded component in the SDDC Manager UI as described below.

With VMware Cloud Foundation 5.1.1 and later, you can add a component license key as described below, or add a solution license key in the vSphere Client. See [Managing vSphere Licenses](#) for information about using a solution license key for VMware ESXi and vCenter Server. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

You first add the new component license key to SDDC Manager. This must be done once per license instance. You then apply the license key to the component on a per workload domain basis.

1. Add a new component license key to the SDDC Manager inventory.
 - a) In the navigation pane, click **Administration** › **Licensing**.
 - b) On the **Licensing** page, click **+ License Key**.
 - c) Select a product from the drop-down menu.
 - d) Enter the license key.
 - e) Enter a description for the license key.
 - f) Click **Add**.
 - g) Repeat for each license key to be added.
2. Update a license key for a workload domain component.
 - a) In the navigation pane, click **Inventory** › **Workload Domains**.
 - b) On the **Workload Domains** page, click the domain you are upgrading.
 - c) On the **Summary** tab, expand the red error banner, and click **Update Licenses**.
 - d) On the **Update Licenses** page, click **Next**.
 - e) Select the products to update and click **Next**.

- f) For each product, select a new license key from the list, and select the entity to which the licensekey should be applied and click **Next**.
- g) On the Review pane, review each license key and click **Submit**.
The new license keys will be applied to the workload domain. Monitor the task in the **Tasks** pane in SDDC Manager.

Upgrade VI Workload Domains to VMware Cloud Foundation 5.2.x

The management domain in your environment must be upgraded before you upgrade VI workload domains. To upgrade to VMware Cloud Foundation 5.2.x, all VI workload domains in your environment must be at VMware Cloud Foundation 4.5 or higher. If your environment is at a version lower than 4.5, you must upgrade the workload domains to 4.5 and then upgrade to 5.2.x.

Within a VI workload domain, components must be upgraded in the following order.

1. NSX.
2. vCenter Server.
3. ESXi.
4. Workload Management on clusters that have vSphere with Tanzu. Workload Management can be upgraded through vCenter Server. See [Updating the vSphere with Tanzu Environment](#).
5. If you suppressed the Enter Maintenance Mode prechecks for ESXi or NSX, delete the following lines from the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file and restart the LCM service:
`lcm.nsxt.suppress.dry.run.emm.check=true`
`lcm.esx.suppress.dry.run.emm.check.failures=true`
6. If you have stretched clusters in your environment, upgrade the vSAN witness host. See [Upgrade vSAN Witness Host for VMware Cloud Foundation](#).

After all upgrades have completed successfully:

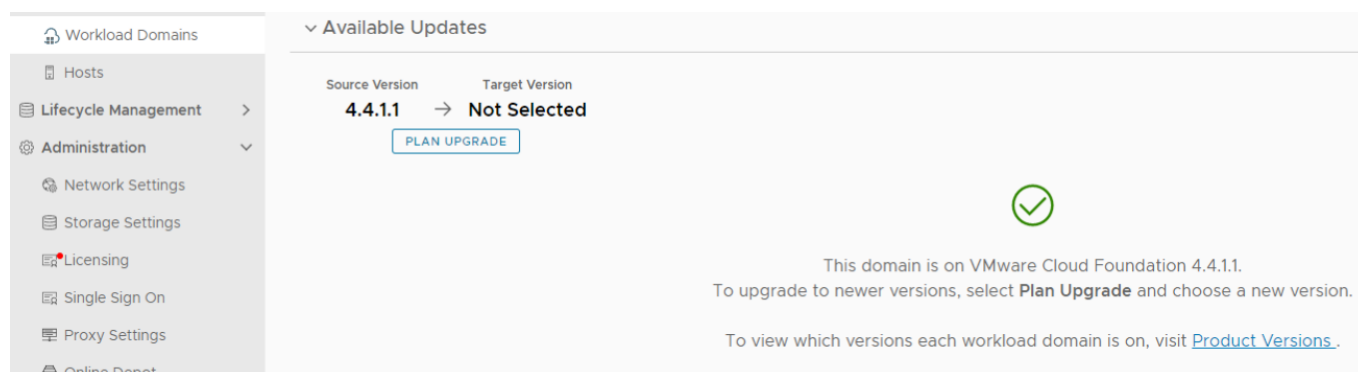
1. Remove the VM snapshots you took before starting the update.
2. Take a backup of the newly installed components.

Plan VI Workload Domain Upgrade

Before proceeding with a VI workload domain upgrade you must first plan the upgrade to your target version.

[Upgrade the Management Domain to VMware Cloud Foundation 5.2.x.](#)

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the **Workload Domains** page, click the workload domain you want to upgrade and click the **Updates** tab.
3. Under **Available Updates**, click **PLAN UPGRADE**.



- On the **Plan Upgrade for VMware Cloud Foundation** screen, select the target version from the drop-down, and click **CONFIRM**.

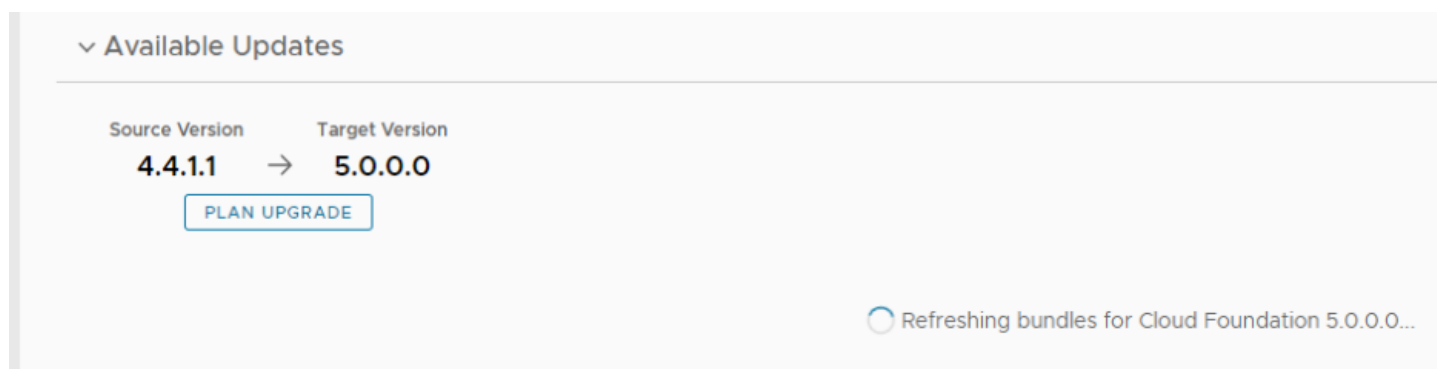
CAUTION

You must upgrade all VI workload domains to VMware Cloud Foundation 5.x. Upgrading to a higher 4.x release once the management domain has been upgraded to 5.x is unsupported.

NOTE

If the target version of VMware Cloud Foundation supports multiple versions of VxRail Manager, the drop-down menu includes separate entries for each combination.

Bundles applicable to the chosen release will be made available to the VI workload domain.



Perform Update Precheck in SDDC Manager

You must perform a precheck in SDDC Manager before applying an update bundle to ensure that your environment is ready for the update.

Bundle-level pre-checks for vCenter are available in VMware Cloud Foundation.

If you silence a vSAN Skyline Health alert in the vSphere Client, SDDC Manager skips the related precheck and indicates which precheck it skipped. Click **RESTORE PRECHECK** to include the silenced precheck. For example: You can also silence failed vSAN prechecks in the SDDC Manager UI by clicking **Silence Precheck**. Silenced prechecks do not trigger warnings or block upgrades.

IMPORTANT

Only silence alerts if you know that they are incorrect. Do not silence alerts for real issues that require remediation.

1. In the navigation pane, click **Inventory** › **Workload Domains**.
2. On the Workload Domains page, click the workload domain where you want to run the precheck.
3. On the domain summary page, click the **Updates** tab.

(The following image is a sample screenshot and may not reflect current product versions.)



NOTE

It is recommended that you Precheck your workload domain prior to performing an upgrade.

4. Click **RUN PRECHECK** to select the components in the workload domain you want to precheck.
 - a) You can select to run a Precheck only on vCenter or the vSphere cluster. All components in the workload domain are selected by default. To perform a precheck on certain components, choose **Custom selection**.



NOTE

For VMware Cloud Foundation on Dell EMC VxRail, you can run prechecks on VxRail Manager.

- b) If there are pending upgrade bundles available, then the "Target Version" dropdown contains "General Upgrade Readiness" and the available VMware Cloud Foundation versions to upgrade to. If there is an available VMware Cloud Foundation upgrade version, there will be extra checks - bundle-level prechecks for hosts, vCenter Server, and so forth. The version specific prechecks will only run prechecks on

components that have available upgrade bundles downloaded.

5. When the precheck begins, a progress message appears indicating the precheck progress and the time when the precheck began.



NOTE

Parallel precheck workflows are supported. If you want to precheck multiple domains, you can repeat steps 1-5 for each of them without waiting for step 5 to finish.

6. Once the Precheck is complete, the report appears. Click through **ALL**, **ERRORS**, **WARNINGS**, and **SILENCED** to filter and browse through the results.



7. To see details for a task, click the expander arrow.

If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **RETRY ALL FAILED RESOURCES** to retry all failed tasks.

8. If the workload domain contains a host that includes pinned VMs, the precheck fails at the Enter Maintenance Mode step. If the host can enter maintenance mode through vCenter Server UI, you can suppress this check for NSX and ESXi in VMware Cloud Foundation by following the steps below.

1. Log in to SDDC Manager by using a Secure Shell (SSH) client with the user name `vcf` and password.
2. Open the `/opt/vmware/vcf/lcm/lcm-app/conf/application-prod.properties` file.
3. Add the following line to the end of the file:

```
lcm.nsxt.suppress.dry.run.emm.check=true
```

```
lcm.esx.suppress.dry.run.emm.check.failures=true
```

- Restart Lifecycle Management by typing the following command in the console window.

```
systemctl restart lcm
```

- After Lifecycle Management is restarted, run the precheck again.

The precheck result is displayed at the top of the Upgrade Precheck Details window. If you click **Exit Details**, the precheck result is displayed at the top of the Precheck section in the Updates tab.

Ensure that the precheck results are green before proceeding. Although a failed precheck will not prevent the upgrade from proceeding, it may cause the update to fail.

Upgrade NSX for VMware Cloud Foundation in a Federated Environment

If NSX Federation is configured between two VMware Cloud Foundation instances, SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must manually upgrade the NSX Global Managers for each instance.

Download NSX Global Manager Upgrade Bundle

SDDC Manager does not manage the lifecycle of the NSX Global Managers. You must download the NSX upgrade bundle manually to upgrade the NSX Global Managers.

- Log in to the Broadcom Support Portal and browse to **My Downloads > VMware NSX**.
- Click the version of NSX to which you are upgrading.
- Locate the **NSX *version* Upgrade Bundle** and verify that the upgrade bundle filename extension ends with `.mub`.

The upgrade bundle filename has the following format `VMware-NSX-upgrade-bundle-versionnumber.buildnumber.mub`.

- Click the download icon to download the upgrade bundle to the system where you access the NSX Global Manager UI.

Upgrade the Upgrade Coordinator for NSX Federation

The upgrade coordinator runs in the NSX Manager. It is a self-contained web application that orchestrates the upgrade process of hosts, NSX Edge cluster, NSX Controller cluster, and the management plane.

The upgrade coordinator guides you through the upgrade sequence. You can track the upgrade process and, if necessary, you can pause and resume the upgrade process from the UI.

- In a web browser, log in to Global Manager for the domain at `https://nsx_gm_vip_fqdn/`.
- Select **System > Upgrade** from the navigation panel.
- Click **Proceed to Upgrade**.
- Navigate to the upgrade bundle `.mub` file you downloaded or paste the download URL link.
 - Click **Browse** to navigate to the location you downloaded the upgrade bundle file.
 - Paste the VMware download portal URL where the upgrade bundle `.mub` file is located.
- Click **Upload**.
When the file is uploaded, the **Begin Upgrade** button appears.
- Click **Begin Upgrade** to upgrade the upgrade coordinator.

NOTE

Upgrade one upgrade coordinator at a time.

7. Read and accept the EULA terms and accept the notification to upgrade the upgrade coordinator..
8. Click **Run Pre-Checks** to verify that all NSX components are ready for upgrade.
The pre-check checks for component connectivity, version compatibility, and component status.
9. Resolve any warning notifications to avoid problems during the upgrade.

Upgrade NSX Global Managers for VMware Cloud Foundation

Manually upgrade the NSX Global Managers when NSX Federation is configured between two VMware Cloud Foundation instances.

Before you can upgrade NSX Global Managers, you must upgrade all VMware Cloud Foundation instances in the NSX Federation, including NSX Local Managers, using SDDC Manager.

1. In a web browser, log in to Global Manager for the domain at https://nsx_gm_vip_fqdn/.
2. Select **System > Upgrade** from the navigation panel.
3. Click **Start** to upgrade the management plane and then click **Accept**.
4. On the Select Upgrade Plan page, select **Plan Your Upgrade** and click **Next**.

The NSX Manager UI, API, and CLI are not accessible until the upgrade finishes and the management plane is restarted.

Upgrade NSX for VMware Cloud Foundation 5.2.x

Upgrade NSX in the management domain and VI workload domains. VMware Cloud Foundation 5.2.1 supports in-place host upgrades for clusters that use vSphere Lifecycle Manager baselines.

Until SDDC Manager is upgraded to version 5.2, you must upgrade NSX in the management domain before you upgrade NSX in a VI workload domain. Once SDDC Manager is at version 5.2 or later, you can upgrade NSX in VI workload domains before or after upgrading NSX in the management domain.

Upgrading NSX involves the following components:

- Upgrade Coordinator
- NSX Edges/Clusters (if deployed)
- Host clusters
- NSX Manager cluster

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates/Patches** tab.

When you upgrade NSX components for a selected VI workload domain, those components are upgraded for all VI workload domains that share the NSX Manager cluster.

3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

NOTE

The NSX precheck runs on all VI workload domains in your environment that share the NSX Manager cluster.

4. For VMware Cloud Foundation 5.2:

- a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for NSX.



- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.

By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.

- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.

By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) If you selected the **Schedule Upgrade** option, specify the date and time for the NSX bundle to be applied and click **Next**.

- f) On the Review page, review your settings and click **Finish**.

If you selected **Upgrade Now**, the NSX upgrade begins and the upgrade components are displayed. The upgrade view displayed here pertains to the workload domain where you applied the bundle. Click the link to the associated workload domains to see the components pertaining to those workload domains. If you selected **Schedule Upgrade**, the upgrade begins at the time and date you specified.

5. For VMware Cloud Foundation 5.2.1:

- a) In the Available Updates section, click the **Configure Update** button.

- b) On the NSX Edge Clusters page, select the NSX Edge clusters you want to upgrade and click **Next**.

By default, all NSX Edge clusters are upgraded. To select specific NSX Edge clusters, select the **Upgrade only NSX Edge clusters** check box and select the **Enable edge selection** option. Then select the NSX Edges you want to upgrade.

- c) On the Host Cluster page, select the host cluster you want to upgrade and click **Next**.

By default, all host clusters across all workload domains are upgraded. If you want to select specific host clusters to upgrade, select **Custom Selection**. Host clusters are upgraded after all Edge clusters have been upgraded.

NOTE

The NSX Manager cluster is upgraded only if you select all host clusters. If you have multiple host clusters and choose to upgrade only some of them, you must go through the NSX upgrade wizard again until all host clusters have been upgraded.

- d) On the Upgrade Options dialog box, select the upgrade optimizations and click **Next**.

By default ESXi hosts are placed into maintenance mode during an upgrade. Starting with VMware Cloud Foundation 5.2.1, in-place upgrades are available for workload domains in which all the clusters use vSphere Lifecycle Manager baselines. If NSX Manager is shared between workload domains, in-place upgrade is only available if all the clusters in all the workload domains that share the NSX Manager use vLCM baselines. If the option is available, you can select **In-place** as the upgrade mode to avoid powering off and placing hosts into maintenance mode before the upgrade.

NOTE

To perform an in-place upgrade, the target NSX version must be the VMware Cloud Foundation 5.2.1 BOM version or later.

By default, Edge clusters and host clusters are upgraded in parallel. You can enable sequential upgrade by selecting the relevant check box.

- e) On the Review page, review your settings and click **Run Precheck**.

The precheck begins. Resolve any issues until the precheck succeeds.

- f) After the precheck succeeds, click **Schedule Update** and select an option.

6. Monitor the upgrade progress. See [monitor-update.dita](#).

If a component upgrade fails, the failure is displayed across all associated workload domains. Resolve the issue and retry the failed task.

When all NSX workload components are upgraded successfully, a message with a green background and check mark is displayed.

Upgrade vCenter Server for VMware Cloud Foundation 5.2.x

The upgrade bundle for VMware vCenter Server is used to upgrade the vCenter Server instances managed by SDDC Manager. Upgrade vCenter Server in the management domain before upgrading vCenter Server in VI workload domains.

- Download the VMware vCenter Server upgrade bundle. See [downloading-vmware-cloud-foundation-bundles.dita](#).
- Take a file-based backup of the vCenter Server appliance before starting the upgrade. See [manually-back-up-vcenter-server.dita](#).

NOTE

After taking a backup, do not make any changes to the vCenter Server inventory or settings until the upgrade completes successfully.

- If your workload domain contains Workload Management (vSphere with Tanzu) enabled clusters, the supported target release depends on the version of Kubernetes (K8s) currently running in the cluster. Older versions of K8s might require a specific upgrade sequence. See [KB 92227](#) for more information.

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

4. Upgrading to VMware Cloud Foundation 5.2:
 - a) In the Available Updates section, click **Update Now** or **Schedule Update** next to the VMware Software Update for vCenter Server.
 - b) Click **Confirm** to confirm that you have taken a file-based backup of the vCenter Server appliance before starting the upgrade.

- c) If you selected **Schedule Update**, click the date and time for the bundle to be applied and click **Schedule**.
- d) If you are upgrading from VMware Cloud Foundation 4.5.x, enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.
-
- e) Review the upgrade settings and click **Finish**.
5. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 5.x:
- a) In the Available Updates section, click **Configure Update**.
- b) Select the upgrade mechanism and click **Next**.

Option	Description
vCenter Reduced Downtime Upgrade	<p>The reduced downtime upgrade process uses a migration-based approach. In this approach, a new vCenter Server Appliance is deployed and the current vCenter data and configuration is copied to it.</p> <p>During the preparation phase of a reduced downtime upgrade, the source vCenter Server Appliance and all resources remain online. The only downtime occurs when the source vCenter Server Appliance is stopped, the configuration is switched over to the target vCenter, and the services are started. The downtime is expected to take approximately 5 minutes under ideal network, CPU, memory, and storage provisioning.</p> <p>NOTE To perform a vCenter Reduced Downtime Upgrade, the target vCenter version must be the VMware Cloud Foundation 5.2.1 BOM version or later.</p>
vCenter Regular Upgrade	During a regular upgrade, the vCenter Server Appliance is offline for the duration of the upgrade.

- c) Select a backup option and click **Next**.
- d) For an RDU update, provide a temporary network to be used only during the upgrade and click **Next**.

Automatic	Automatically assign network information.
Static	Enter an IP address, subnet mask, and gateway. The IP address must be in the management subnet.

- e) Schedule the update and click **Next**.

For vCenter Reduced Downtime Upgrade	Select scheduling options for the preparation and switchover phases of the upgrade.
--------------------------------------	---

Table continued on next page

Continued from previous page

	<p>NOTE If you are scheduling the switchover phase, you must allow a minimum of 4 hours between the start of preparation and the start of switchover.</p>
For vCenter Regular Upgrade	Select an Upgrade Now or Schedule Update .

- f) Review the upgrade settings and click **Finish**.
6. Upgrading to VMware Cloud Foundation 5.2.1 from VMware Cloud Foundation 4.5.x:
 - a) In the Available Updates section, click **Configure Update**.
 - b) Enter the details for the temporary network to be used only during the upgrade. The IP address must be in the management subnet.
 - c) Select a backup option and click **Next**.
 - d) Schedule the update and click **Next**.
 - e) Review the upgrade settings and click **Finish**.
7. Monitor the upgrade progress. See [monitor-update.dita](#).
8. After the upgrade is complete, remove the old vCenter Server appliance (if applicable).

NOTE

Removing the old vCenter is only required for major upgrades. If you performed a vCenter RDU patch upgrade, the old vCenter is automatically removed after a successful upgrade.

If the upgrade fails, resolve the issue and retry the failed task. If you cannot resolve the issue, restore vCenter Server using the file-based backup. See [restore-vcenter-server.dita](#). vCenter RDU upgrades perform automatic rollback if the upgrade fails.

Once the upgrade successfully completes, use the vSphere Client to change the vSphere DRS Automation Level setting back to the original value (before you took a file-based backup) for each vSphere cluster that is managed by the vCenter Server. See [KB 87631](#) for information about using VMware PowerCLI to change the vSphere DRS Automation Level.

Upgrade VxRail Manager and ESXi Hosts for VMware Cloud Foundation

Use the VxRail upgrade bundle to upgrade VxRail Manager and the ESXi hosts in the workload domain. Upgrade the management domain first and then VI workload domains.

- Validate that the ESXi passwords are valid.
- Download the VxRail upgrade bundle. See [downloading-vmware-cloud-foundation-bundles.dita](#).
- Ensure that the domain for which you want to perform cluster-level upgrade does not have any hosts or clusters in an error state. Resolve the error state or remove the hosts and clusters with errors before proceeding.

By default, the upgrade process upgrades the ESXi hosts in all clusters in a workload domain in parallel. If you have multiple clusters in the management domain or in a VI workload domain, you can select the clusters to upgrade. You can also choose to upgrade the clusters in parallel or sequentially.

If you are using external (non-vSAN) storage, the following procedure updates the ESXi hosts attached to the external storage. However, updating and patching the storage software and drivers is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability after an ESXi upgrade, consult the vSphere HCL and your storage vendor.

1. Navigate to the **Updates/Patches** tab of the workload domain.
2. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with the upgrade.

3. In the Available Updates section, select the target release.
4. Click **Upgrade Now** or **Schedule Update**.
If you selected **Schedule Update**, specify the date and time for the bundle to be applied.
5. Select the clusters to upgrade and click **Next**.
The default setting is to upgrade all clusters. To upgrade specific clusters, click **Enable cluster-level selection** and select the clusters to upgrade.
6. Click **Next**.
7. Select the upgrade options and click **Finish**.
By default, the selected clusters are upgraded in parallel. If you selected more than five clusters to be upgraded, the first five are upgraded in parallel and the remaining clusters are upgraded sequentially. To upgrade all selected clusters sequentially, select **Enable sequential cluster upgrade**.

Click **Enable Quick Boot** if desired. Quick Boot for ESXi hosts is an option that allows Update Manager to reduce the upgrade time by skipping the physical reboot of the host.
8. Monitor the upgrade progress. See [monitor-update.dita](#).

Upgrade the vSAN Disk Format for vSAN clusters. The disk format upgrade is optional. Your vSAN cluster continues to run smoothly if you use a previous disk format version. For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of vSAN. See [upgrade-vsan-on-disk-format-versions.dita](#).

Upgrade vSAN Witness Host for VMware Cloud Foundation

If your VMware Cloud Foundation environment contains stretched clusters, update and remediate the vSAN witness host.

Download the ESXi ISO that matches the version listed in the the Bill of Materials (BOM) section of the *VMware Cloud Foundation Release Notes*.

1. In a web browser, log in to vCenter Server at https://vcenter_server_fqdn/ui.
2. Upload the ESXi ISO image file to vSphere Lifecycle Manager.
 - a) Click **Menu** > **Lifecycle Manager**.
 - b) Click the **Imported ISOs** tab.
 - c) Click **Import ISO** and then click **Browse**.
 - d) Navigate to the ESXi ISO file you downloaded and click **Open**.
 - e) After the file is imported, click **Close**.
3. Create a baseline for the ESXi image.
 - a) On the Imported ISOs tab, select the ISO file that you imported, and click **New baseline**.
 - b) Enter a name for the baseline and specify the **Content Type** as Upgrade.
 - c) Click **Next**.
 - d) Select the ISO file you had imported and click **Next**.
 - e) Review the details and click **Finish**.
4. Attach the baseline to the vSAN witness host.
 - a) Click **Menu** > **Hosts and Clusters**.
 - b) In the Inventory panel, click **vCenter** > **Datacenter**.
 - c) Select the vSAN witness host and click the **Updates** tab.
 - d) Under Attached Baselines, click **Attach** > **Attach Baseline or Baseline Group**.
 - e) Select the baseline that you had created in step 3 and click **Attach**.
 - f) Click **Check Compliance**.

After the compliance check is completed, the **Status** column for the baseline is displayed as Non-Compliant.

5. Remediate the vSAN witness host and update the ESXi hosts that it contains.
 - a) Right-click the vSAN witness and click **Maintenance Mode › Enter Maintenance Mode**.
 - b) Click **OK**.
 - c) Click the **Updates** tab.
 - d) Select the baseline that you had created in step 3 and click **Remediate**.
 - e) In the End user license agreement dialog box, select the check box and click **OK**.
 - f) In the Remediate dialog box, select the vSAN witness host, and click **Remediate**.

The remediation process might take several minutes. After the remediation is completed, the **Status** column for the baseline is displayed as Compliant.

- g) Right-click the vSAN witness host and click **Maintenance Mode › Exit Maintenance Mode**.
- h) Click **OK**.

Upgrade vSphere Distributed Switch versions

[Optional] Upgrade the distributed switch to take advantage of features that are available only in the later versions.

ESXi and vCenter Upgrades are completed.

1. On the vSphere Client Home page, click **Networking** and navigate to the distributed switch.
2. Right-click the distributed switch and select **Upgrade › Upgrade Distributed Switch**
3. Select the vSphere Distributed Switch version that you want to upgrade the switch to and click **Next**

The vSphere Distributed Switch is successfully upgraded.

Upgrade vSAN on-disk format versions

[Optional] Upgrade the vSAN on-disk format version to take advantage of features that are available only in the later versions.

- ESXi and vCenter Upgrades are completed
- Verify that the disks are in a healthy state. Navigate to the Disk Management page to verify the object status.
- Verify that your hosts are not in maintenance mode. When upgrading the disk format, do not place the hosts in maintenance mode.
- Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster. For information about vSAN resynchronization, see vSphere Monitoring and Performance
- The upgrade may cause temporary resynchronization traffic and use additional space by moving data or rebuilding object components to a new data structure.

1. Navigate to the vSAN cluster.
2. Click the **Configure** tab.
3. Under **vSAN**, select **Disk Management**.
4. Click **Pre-check Upgrade**. The upgrade pre-check analyzes the cluster to uncover any issues that might prevent a successful upgrade. Some of the items checked are host status, disk status, network status, and object status. Upgrade issues are displayed in the **Disk pre-check status text** box.
5. Click **Upgrade**.
6. Click **Yes** on the Upgrade dialog box to perform the upgrade of the on-disk format.

vSAN successfully upgrades the on-disk format. The On-disk Format Version column displays the disk format version of storage devices in the cluster

Update License Keys for a Workload Domain

If upgrading from a VMware Cloud Foundation version prior to 5.0, you need to update your license keys to support vSAN 8.x and vSphere 8.x.

You need a new license key for vSAN 8.x and vSphere 8.x. Prior to VMware Cloud Foundation 5.1.1, you must add and update the component license key for each upgraded component in the SDDC Manager UI as described below.

With VMware Cloud Foundation 5.1.1 and later, you can add a component license key as described below, or add a solution license key in the vSphere Client. See [Managing vSphere Licenses](#) for information about using a solution license key for VMware ESXi and vCenter Server. If you are using a solution license key, you must also add a VMware vSAN license key for vSAN clusters. See [Configure License Settings for a vSAN Cluster](#).

You first add the new component license key to SDDC Manager. This must be done once per license instance. You then apply the license key to the component on a per workload domain basis.

1. Add a new component license key to the SDDC Manager inventory.
 - a) In the navigation pane, click **Administration** > **Licensing**.
 - b) On the **Licensing** page, click **+ License Key**.
 - c) Select a product from the drop-down menu.
 - d) Enter the license key.
 - e) Enter a description for the license key.
 - f) Click **Add**.
 - g) Repeat for each license key to be added.
2. Update a license key for a workload domain component.
 - a) In the navigation pane, click **Inventory** > **Workload Domains**.
 - b) On the **Workload Domains** page, click the domain you are upgrading.
 - c) On the **Summary** tab, expand the red error banner, and click **Update Licenses**.
 - d) On the **Update Licenses** page, click **Next**.
 - e) Select the products to update and click **Next**.
 - f) For each product, select a new license key from the list, and select the entity to which the licensekey should be applied and click **Next**.
 - g) On the Review pane, review each license key and click **Submit**.
The new license keys will be applied to the workload domain. Monitor the task in the **Tasks** pane in SDDC Manager.

Independent SDDC Manager Upgrade using the SDDC Manager UI

Once SDDC Manager is upgraded to 5.2 or later, new functionality is introduced that allows you to get the latest SDDC Manager features and security fixes without having to upgrade the entire VMware Cloud Foundation BOM. An independent SDDC Manager release includes a fourth digit in its version number, for example SDDC Manager 5.2.0.1.

- Download the SDDC Manager bundle. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- SDDC Manager must be version 5.2 or later.

You can upgrade SDDC Manager without upgrading the full VCF BOM when:

- The target version of SDDC Manager is compatible with all the BOM product versions running in your current environment (management and workload domains).
- There is a supported upgrade path from your current SDDC Manager version to the target SDDC Manager version.

NOTE

You can use the SDDC Manager upgrade functionality to upgrade SDDC Manager even when the target version of SDDC Manager is part of a full VCF BOM release, as long as it is compatible.

Updating SDDC Manager without upgrading the full VCF BOM, does not change the version of the management domain.

1. In the navigation pane, browse to **Lifecycle Management > SDDC Manager**.
The UI displays available SDDC Manager updates that are either SDDC Manager only updates or SDDC Manager updates that are part of a full VCF BOM update.
2. Review and address any compatibility warnings.
3. Click **Run Precheck**.
Resolve any precheck issues before proceeding.
4. Schedule the update to run now or at a specific time and click **Start Update**.
When the update completes successfully, you are logged out of the SDDC Manager UI and must log in again.

Flexible BOM Upgrade in VMware Cloud Foundation

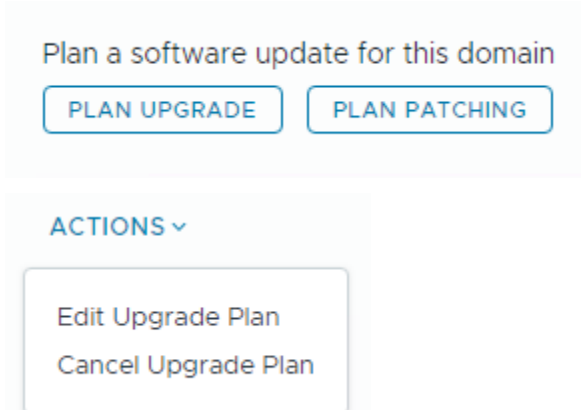
Once SDDC Manager is upgraded to version 5.2 or later, new functionality is introduced to the upgrade planner that allows you to select specific target versions for each VMware Cloud Foundation component you want to upgrade.

- Download the bundles for the target versions of each VCF component. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- SDDC Manager must be version 5.2 or later.

You can use the upgrade planner to select any supported version for each of the VMware Cloud Foundation BOM components. This includes async patch versions as well as VCF BOM versions.

To plan an upgrade when SDDC Manager does not have internet access, see [Offline Download of Flexible BOM Upgrade Bundles](#).

1. In the navigation pane, click **Inventory > Workload Domains**.
2. On the Workload Domains page, click the domain you are upgrading and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.
Resolve any issues before proceeding with an upgrade.
4. In the Available Updates section, click **Plan Upgrade** create a new upgrade plan or select **Edit Upgrade Plan** from the **Actions** menu to modify an upgrade plan.



5. Select the target version of VMware Cloud Foundation and VxRail Manager from the drop-down menu and click **Next**.

Plan Upgrade

Select a version of VMware Cloud Foundation to upgrade this workload domain.

To activate or deactivate compatibility checks on the target version, see KB [90074](#).

Target Version: Select VMware Cloud Foundation with VxRail Manager Version ▾

NEXT CANCEL

Change Summary

The following product upgrades will be planned out. The target versions of individual products can be customized if necessary.

VALIDATE SELECTION **RESET TO DEFAULTS**

Software Component	Current Version	Current Build	Target Version	Target Build
SDDC Manager	5.2.0.0	24108943	<u>5.2.0.0-24108943</u> ▾	24108943
VMware NSX	4.1.2.0.0	22305534	<u>4.1.2.1.0-22667789</u> ▾	22667789
VMware vCenter Server Appliance	8.0.2.00000	23718066	<u>8.0.2.00100-22617221</u>	22617221
VxRail Manager	8.0.000	27667632	<u>8.0.200-28314729</u> ▾	28314729

CONFIRM BACK CANCEL

- After validation succeeds, click **Confirm**.
- Review the update sequence based on your target version selections and click **Done**.
- In the Available Updates screen, click **Schedule Update** or **Update Now** to update the first component. Continue to update the VCF BOM components until they are all updated.

NOTE

If SDDC Manager does not have internet access, you need to perform additional steps before you can start updating. See [Offline Download of Flexible BOM Upgrade Bundles](#).

Patching the Management and Workload Domains

Once SDDC Manager is upgraded to 5.2 or later, a new option for patching VMware Cloud Foundation components is available in the SDDC Manager UI.

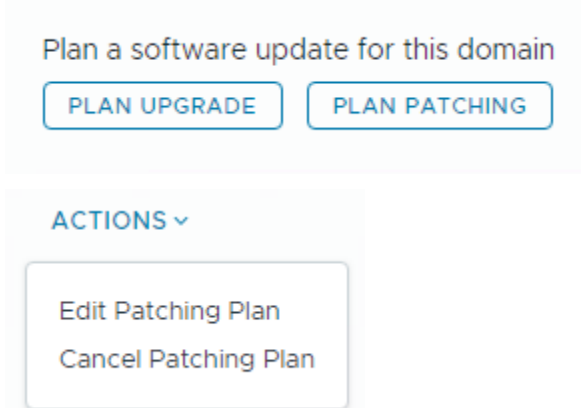
- Download the async patch bundles. See [Downloading VMware Cloud Foundation Upgrade Bundles](#).
- SDDC Manager must be version 5.2 or later. See [Apply the VMware Cloud Foundation 5.2.x Upgrade Bundle](#).

The patch planner provides the ability to apply async patches to workload domain components. If you are connected to the online depot, async patches are available in the patch planner. If you do not have access to the online depot, use the Bundle Transfer Utility to download async patches and add them to an offline depot or upload them directly to SDDC Manager.

1. In the navigation pane, click **Inventory** > **Workload Domains**.
2. On the Workload Domains page, click the domain you are patching and then click the **Updates** tab.
3. Click **Precheck** to run the upgrade precheck.

Resolve any issues before proceeding with an upgrade.

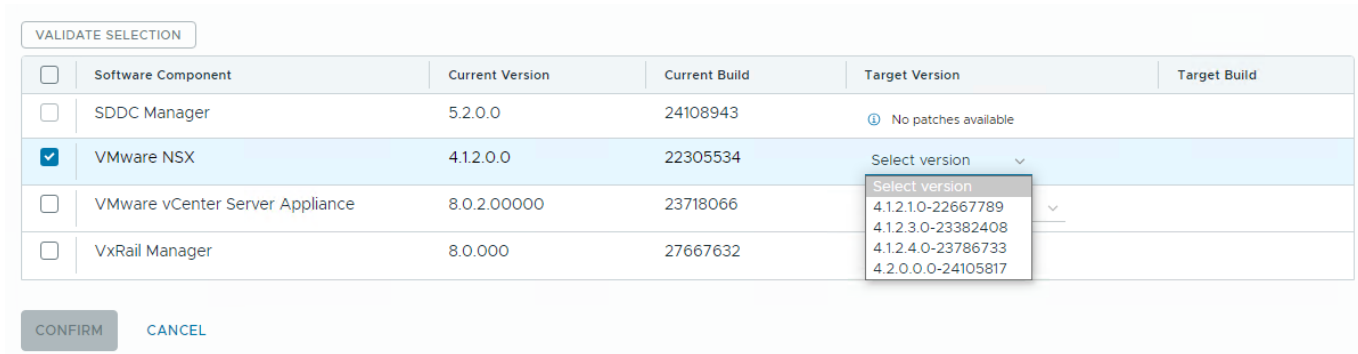
4. In the Available Updates section, click **Plan Patching** create a new patching plan or select **Edit Patching Plan** from the **Actions** menu to modify a patching plan.



NOTE

You cannot plan patching if you have an existing upgrade plan. Cancel the upgrade plan to create a patching plan.

5. Select the components to patch and the target versions and then click **Validate Selection**.




NOTE


When you select a target vCenter version, the UI indicates which versions support vCenter Reduced Downtime Upgrade (RDU).

6. After validation succeeds, click **Confirm**.
7. Review the update sequence based on your target version selections and click **Done**.

Plan Patching

 Patching plan confirmed. See your updates summary below.

Update sequence preview

Step 1  NSX_T_MANAGER Update Bundle 4.2.0.0.0

The upgrade bundle for VMware NSX Data Center 4.2.0.0. Customers are strongly encouraged to run the NSX Upgrade Evaluation Tool. For more information, see <https://docs.vmware.com/en/VMware-NSX/4.2/rn/vmware-nsxt-data-center-42-release-notes/index.html>.

DONE

[BACK AND EDIT PATCHES](#)

- In the Available Updates screen, click **Schedule Update** or **Update Now** to update the first component. Continue to update the VCF BOM components until they are all updated.

Troubleshooting for Upgrading VMware Cloud Foundation

A library of troubleshooting processes that may be referenced during the VMware Cloud Foundation upgrade as appropriate.

SDDC Manager Troubleshooting

A library of SDDC Manager troubleshooting processes that may be referenced during upgrade as appropriate.

On-demand pre-checks for vCenter bundle might fail

The bundle pre-check failure can occur in a specific scenario. When SDDC Manager is upgraded to VMware Cloud Foundation 5.0.0.x from 4.5.x, and BOM components are not upgraded to VMware Cloud Foundation 5.0.0.x and Customer downloads the bundles for VMware Cloud Foundation 5.1.0.0 and runs the pre-check by selecting target version as 5.1.0.0.

The format of the vCenter Server bundle is modified starting from VMware Cloud Foundation 5.1. The new bundle is a unified bundle that bundles both the .iso and .zip files for the Target vCenter Server build. This unified bundle can be used for both major and minor vCenter Server upgrades. The SDDC Manager needs to be at least at the 5.1 version to understand the new format and run the prechecks. As VMware Cloud Foundation 5.0.0.0 does not understand the format, the bundle pre-check will fail.

Error Message: Upgrade Bundle Validation

- Upgrade the SDDC Manager to VMware Cloud Foundation 5.1.0.0 and run the on-demand prechecks for vCenter Server in VMware Cloud Foundation 5.1.0.0.

<https://kb.vmware.com/s/article/94862>

SDDC Manager bundle pre-check failure when upgrading to VMware Cloud Foundation 5.1

SDDC Manager Pre-check fails

SDDC Manager Pre-check "Upgrade Bundle Download Status" fails with an error

- "Could not find bundle for SDDC_MANAGER upgrade to version 5.1.0.0-<build_number>".

From VMware Cloud Foundation 5.1 onwards, we are deprecating the Config Drift bundle. However, the previously released VCF versions expect that a config drift bundle will be applied as part of a target release and hence indicate this as a pre-check failure.

This pre-check failure can be ignored for VCF 5.1+, and it is safe to proceed with the upgrade despite this bundle pre-check failure.

Extra RPM packages on SDDC Manager may cause upgrade failure

SDDC Manager upgrade may fail if some RPMs on the current SDDC Manager are incompatible with those on the upgraded SDDC Manager. In `/var/log/vmware/capengine/cap-update/install-*`,

You may see a message like:

- package `systemd-udev-247.13-4.ph4.x86_64` requires `libsystemd-shared-247.so()(64bit)`, but none of the providers can be installed.
 - package `systemd-247.13-4.ph4.x86_64` requires `libcrypto.so.3()(64bit)`, but none of the providers can be installed.
 - package `rpm-4.16.1.3-17.ph4.x86_64` requires `libcrypto.so.3()(64bit)`, but none of the providers can be installed
- RPMs may have been left behind by previous upgrades or greenfield deployments, or a user has implicitly or explicitly installed RPMs that prevent the upgrade

The workaround is to uninstall RPMs that are causing this upgrade conflict manually.

False warning for missing compatibility data in plan upgrade wizard

When no compatibility data is missing, an incorrect warning message is populated

A warning message with an empty product list in the plan upgrade wizard appears

- "Unable to verify the compatibility for the following product versions. Please check the product documentation before proceeding to upgrade:"

Users can ignore the warning and is not blocked.

Updating licenses for a WLD shows insufficient license error

When the 'Update Licenses' operation is performed for a Workload Domain, in certain cases, the incorrect quantity of licenses is shown in the 'Available quantity' field

This is due to a miscalculation in the no. of available licenses. Along with the incorrect quantity, an error alert might also be displayed saying,

- 'License key has insufficient license.'

A miscalculation in the code for the number of available licenses is causing the error alert to appear.

The users can simply choose to ignore the incorrect license count in the 'Available quantity' field when assigning the license. Also, the error alert should be ignored as it does not prohibit the user from moving forward. Users can proceed with the addition of a license even with the error alert. If there are sufficient licenses available, the operation will succeed.

vCenter Troubleshooting

A library of vCenter troubleshooting processes that may be referenced during upgrade as appropriate.

vCenter Server Upgrade Failed Due to Reuse of Temporary IP Address

vCenter Server Upgrade Failed Due to Reuse of Temporary IP Address with error "Cannot run the revert networking command. `revert_networking.py` doesn't exist on target VC" or "VC upgrade is failing during Install-"target vc upgrade precheck stage failing"

Reuse of temporary IP address causes an arp cache issue. Reset the arp cache on the management domain vCenter Server.

Customers who have fewer Temporary IP Addresses than vCenter Servers that are conducting a parallel upgrade have the highest likelihood of impact.

1. SSH to the management domain vCenter Server as root.
2. Run the following

```
ip -s -s neigh flush all
```

Shutdown and Startup of VMware Cloud Foundation

Shutting down VMware Cloud Foundation, for example, during hardware or power maintenance, and then starting it up must be done in a way that prevents data loss or appliance malfunction, and supports collection of troubleshooting data.

You follow a strict order and steps for shutdown and startup of the VMware Cloud Foundation management components.

Shutting Down VMware Cloud Foundation

To avoid data loss and maintain the SDDC components operational, you follow a specific order when shutting down the management virtual machines in VMware Cloud Foundation.

- Verify that you have complete backups of all management components.
- Verify that the management virtual machines are not running on snapshots.
- If a vSphere Storage APIs for Data Protection (VADP) based backup solution is running on the management clusters, verify that the solution is properly shut down by following the vendor guidance.
- To reduce the startup time before you shut down the management virtual machines, migrate the VMware vCenter Server® instance for the management domain to the first VMware ESXi™ host in the default management cluster in the management domain.

You shut down the customer workloads and the management components for the VI workload domains before you shut down the components for the management domain.

If the VMware NSX® Manager™ cluster and VMware NSX® Edge™ cluster are shared with other VI workload domains, shut down the NSX Manager and NSX Edge clusters as part of the shutdown of the first VI workload domain.

Starting Up VMware Cloud Foundation

To maintain the components integration and avoid operation faults, you follow a specified order to start up the management virtual machines in VMware Cloud Foundation.

- Verify that external services such as Active Directory, DNS, NTP, SMTP, and FTP or SFTP are available.
- If a vSphere Storage APIs for Data Protection (VADP) based backup solution is deployed on the default management cluster, verify that the solution is properly started and operational according to the vendor guidance.

You start the management components for the management domain first. Then, you start the management components for the VI workload domains and the customer workloads.

If the NSX Manager cluster and NSX Edge cluster are shared with other VI workload domains, start the other VI workload domains first. Start up NSX Manager and NSX Edge nodes as part of the startup of the last workload domain.

Instance Recovery Guide

Recovering a VMware Cloud Foundation system by performing a complete reconstruction from a backup.

This document provides detailed instructions on recovering an entire VMware Cloud Foundation system, including the management domain and VI workload domains, where you must recover all components.

Example Failure Scenarios

The cases when you must recover all components in a VMware Cloud Foundation instance might be one of the following:

- Complete site failure
- Recovery from a malware or ransomware attack
- Catastrophic logical corruption

Intended Audience

VMware Cloud Foundation Instance Recovery Guide is intended for cloud architects, cloud administrators, and cloud operators who are familiar with and want to recover a VMware Cloud Foundation system that has experienced a significant failure.

Related VMware Cloud Foundation Documentation

In addition to this documentation, the following publications for the VMware Cloud Foundation version in your environment must be available during the recovery process:

- VMware Cloud Foundation Deployment Guide
- VMware Cloud Foundation Administration Guide

You can open these documents from the [VMware Cloud Foundation documentation](#) main page.

Verifications and Remediations

After you recover the default cluster of the management domain or a cluster in a VI workload domain, verify the health of the recovered system and remediate any issues.

Check SDDC Manager Health

After rebuilding a cluster, you should perform a health check using the SoS utility.

1. SSH into the SDDC Manager VM as `vcf`.
2. Enter the following commands:

```
su -  
cd /opt/vmware/sddc-support  
./sos --health-check --domain-name ALL --skip-cert-check
```

3. Remediate any issues reported by the health check.

Delete the Temporary vCenter Server Instance

You can now delete the temporary vCenter Server deployed at partial bring-up.

1. Power off and delete the temporary vCenter Server.

Perform SDDC Manager Pre-Checks

Perform an upgrade pre-check on each workload domain to ensure the system is in a healthy state.

1. Log in to the SDDC Manager UI as **administrator@vsphere.local**.
2. In the navigation pane, click **Inventory › Workload Domains**.
3. On the **Workload Domains** page, click the workload domain where you want to run the pre-check operation.
4. On the **Summary** tab, click the **Updates** tab.
5. Click **Run Precheck** to validate that the environment is ready to be upgraded.
6. From the **Target Version** drop-down menu, select **General Upgrade Readiness**.
7. Configure the precheck to run on the entire workload domain.
8. Click **Run Precheck**.
9. Click **View Status** to see detailed tasks and their status.
10. To see details for a task, click the Expand arrow.

If a precheck task failed, fix the issue, and click **Retry Precheck** to run the task again. You can also click **Precheck Failed Resources** to retry all failed tasks.

Introducing Security and Compliance for VMware Cloud Foundation 5.2

The *Introducing Security and Compliance for VMware Cloud Foundation* document provides general guidance for organizations that are considering VMware solutions to help them address on-premise compliance requirements. This document is a building block of the *Compliance Kit for VMware Cloud Foundation*.

WARNING

This document is intended to provide general guidance for organizations that are considering VMware solutions to help them address on-premise compliance requirements. The information contained in this document is for educational and informational purposes only. This document is not intended to provide regulatory advice. VMware makes no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained in this document. Organizations should engage appropriate legal, business, technical, and audit expertise for review of regulatory compliance requirements.

Intended Audience

Introducing Security and Compliance for VMware Cloud Foundation is intended for cloud architects, infrastructure administrators, and cloud administrators. Familiarity with VMware software is required. This guide introduces security and compliance as it relates to the VMware Cloud Foundation.

Required VMware Software

The *Introducing Security and Compliance for VMware Cloud Foundation* document builds on top of VMware Cloud Foundation and is specific to the standard architecture model of VMware Cloud Foundation.

The products included in this compliance kit cover selected products from the VMware Cloud Foundation 5.2 bill-of-materials:

- VMware ESXi™
- VMware vCenter Server®
- VMware vSAN™
- VMware NSX™
- VMware Cloud Foundation™ SDDC Manager™

See *VMware Cloud Foundation Release Notes* for more information about supported product versions.

Security by Design

Security and compliance guidance includes both default configurations in the VMware Cloud Foundation and non-default configurations that can be implemented post-deployment.

The *Compliance Kit for VMware Cloud Foundation* views configurations from two personas. System administrators and implementation teams for VMware Cloud Foundation use the *Security and Compliance Configuration for VMware Cloud Foundation* to assess and implement non-default configurations. Default configurations that address compliance are not subject of the configuration guide because they do not require additional configuration. In some cases, default configurations must be evaluated to ensure the default parameter aligns with the policy and procedures of your organization. Guidance for auditors who evaluate a VMware Cloud Foundation environment can use the *VMware Cloud Foundation Audit Guide Appendix* to evaluate both default and non-default configurations.

Default configurations

Security configurations based on compliance requirements that are configured by default in VMware Cloud Foundation. According to the different regulatory requirements, the parameter values might require changes, but by secure design these configurations are included in the current implementation.

Non-default configurations

Additional input by the organization is required to identify, select, and set configurations based on a target regulation.

Security Architecture

Security in VMware Cloud Foundation is evaluated with a clear objective to balance best practices with usability and performance.

For VMware Cloud Foundation implementations, post-deployment, security must be handed over to a dedicated team to augment and monitor the security posture. Attack vectors and compliance guidelines are constantly evolving so the information provided is often used to establish a baseline, not an absolute, or complete picture.

NIST 800-53 Revision 5, risk rating Moderate, forms the security baseline used to evaluate VMware Cloud Foundation. NIST 800-53 is the baseline because of its vast array of controls and because it is often used by other regulations as part of their reference framework.

NIST is a risk-based framework, which requires each organization to assess their own risk posture and identify applicable controls. The *Compliance Kit for VMware Cloud Foundation* does not remove this step. The VMware Cloud Foundation security design and compliance mappings inform the reader of both design decisions and security configurations.

The VMware Cloud Foundation security design is not enough on its own. Each organization must have a series of supporting security architecture, technology, processes, and people to evaluate. Applications, workload domains, software-defined networking topology, customer data, privacy, and myriad other factors must be evaluated as part of the overall security architecture.

Super users of the system inherit various technologies and typically work with security specialists to implement controls effectively. VMware Cloud Foundation has evaluated many design decisions that are incorporated with the overall design as outlined by VMware Validated Design architecture guides.

Subsequent deployments benefit from post-implementation security health checks to enhance the organizations security posture as it relates to the VMware Validated Design used in conjunction with VMware Cloud Foundation.

Security Principles

Across all regulations or standards, security principles dictate the mindset for applying security controls in VMware Cloud Foundation.

The security concepts are treated as guiding principles to develop a secure VMware Cloud Foundation environment that leverages capabilities available across all products. These principles do not only result in the configurations identified in this guidance but are also inherent in product capabilities. Organizations that leverage these guidelines can expand these capabilities across the Software-Defined Data Center to include people, process, and technology controls. Each organization must tailor these principles and prioritize how they approach them.

Separation of duties

- Assign roles to users to separate conflicts of duty
- Roles can be customized and further tailored as needed.
- Restrict the use of super users
- Create service accounts where possible
- Create separate accounts for system-to-system communication
- Separate production from development environments
- Evaluate access to create, edit, or delete permissions
- Assign only read-only access where possible

Least privilege

- Deactivate unused services
- Do not grant or retain permissions longer than needed

Confidentiality - Integrity - Availability (CIA)

- Protect the data and the assets used to access it
- Confidentiality applies to the authorization to access the data
- Integrity applies to the authorization to modify the data
- Availability applies to the accessibility to access the data

Defense in depth

- Do not allow lateral movement
- Isolate environments
- Patch systems
- Implement layered security

Zero trust

- Implicit access denial regardless of origin
- Treat internal network as a potential threat vector
- Access is restricted via a trust broker
- Applications are hidden from discovery

Secure Software Development Life-Cycle (SDLC)

- VMware performs static code analysis
- VMware performs penetration testing
- VMware performs vulnerability scan
- Align development with VMware internal vSECR software development security guidelines/procedures

Data in transit protection

- Encryption of virtual machines during migration between hosts
- Use of encrypted mechanism when a super user is interacting with server consoles

Data at rest protection

- Encryption of virtual machines while powered off (at rest)

Trusted Computing Base (TCB)

- Architecture view that brings together the collection of all the hardware, software, and firmware components (including the security kernel and reference monitor)
- Brings a unified security policy and baseline consistent across various layers, abstractions, and detailed components to meet security requirements.

Governance, Risk, and Compliance and Mapping

This guidance describes the security configurations that can support Governance, Risk, and Compliance (GRC) considerations. Due to the variety of compliance standards and different organizational business needs, due care should be taken to identify and map VMware Cloud Foundation configurations against a targeted regulation.

Where possible, examples of audit artifacts are included as evidence in the *VMware Cloud Foundation Audit Guide Appendix*, focused on compliance and producing evidence to meet controls. To map configurations across regulatory standards, we use a third-party tool produced by the Unified Compliance Framework (UCF). This removes a subjective, manual control cross-walk approach and replaces it with a repeatable and data driven methodology. The crosswalk or reference across regulatory standards is not a mapping matrix, but instead utilizes the UCF as a shared

library of controls tied to the underlying citation text within each standard. This removes the subjective mapping and replaces it with a programmatic, software-driven mapping engine.

In some cases, the regulation may be too generic or too vague, which can reduce the mapping efficacy. In these cases, an additional review is performed to isolate new citation text and then included in the engine through the corresponding and newly identified UCF control. No mapping is provided with an accompany UCF control and accompanying citation text for each regulation. If no mapping is identified, the mapping uses `VMware Best Practice` text to clarify that mapping was not found but to keep up with the security principles, the configuration is recommended.

The compliance mapping is a subject of expansion, as more security controls are evaluated, including additional compliance domains and regulations.

Control Definition

Controls are designed to mitigate risk. These are derived by using a Risk Framework, such as the *Guide for Applying the Risk Management Framework to Federal Information Systems* published by NIST, publication number 800-37. NIST 800-53 R5 control catalog is used to develop a baseline of controls compared to the software-defined data center technical and security configurations. These security configurations must be evaluated and considered against the risk management framework used by your organization. Other frameworks such as ISO27001 can be coupled with its Annex A, ISO27002, or ISO27005 to evaluate controls to mitigate risk.

Cybersecurity Considerations

It is the responsibility of each security, compliance, and audit teams in your organization to verify that configurations meet their compliance requirements. The attack vectors and compliance guidelines are constantly evolving, which requires constant monitoring and risk management processes.

Business Impact Assessment

Measuring risk and evaluating scope may require performing a business impact assessment. This analysis can inform IT security and audit professionals the areas of the Software-Defined Data Center that require more controls, tightened access restrictions, micro-segmentation, enhanced disaster recovery, and additional monitoring.

Compliance Kit for VMware Cloud Foundation

The compliance kits is a solution that builds on top of VMware Cloud Foundation and leverages security fundamentals. The kit address the top ten most frequently requested compliance standards, regulations, and frameworks.

The compliance kit is designed and validated to tailor security configurations without impacting the ability of VMware Cloud Foundation to meet its design objectives. The kit can assist organizations to secure information systems in a compliance context.

This guidance has been validated and tested against certain product versions. Changes between subsequent releases of VMware Cloud Foundation are designed for stability and optimal upgrade experience. Guidance provided by the *Compliance Kit for VMware Cloud Foundation* is for a specific VMware Cloud Foundation release, but can still be used until a subsequent kit release is available. This guidance is not backward-compatible and must not be implemented for separate product components.

You can directly download the [VMware Cloud Foundation 5.2 Compliance Kit](#).

Compliance Kit for VMware Cloud Foundation Structure

The compliance kit consists of documents specific to the standard architecture model of VMware Cloud Foundation.

Document Name	Document Description	Intended Audience
<i>Security and Compliance Configuration for VMware Cloud Foundation</i>	Non-default configurations can be performed post deployment of VMware Cloud Foundation for Standard Architecture.	<ul style="list-style-type: none"> • System Integrator • Cloud Administrator • Infrastructure Administrator
<i>VMware Cloud Foundation Audit Guide Appendix</i>	Includes audit procedures for auditors examining an environment for compliance readiness.	<ul style="list-style-type: none"> • System Integrator • Cloud Administrator • Security Professional • Auditor

The compliance kit is designed to work holistically. Each document supports the overall blueprint and builds trust across multiple persona that may interact with the life cycle of a system operating within a compliance context: architect, system administrator, system integrator, security professional, and auditor.

Introducing Security and Compliance for VMware Cloud Foundation outlines security and compliance concepts used in the development of the VMware Cloud Foundation, Compliance Kit. For example, considerations such as governance, risk, and compliance, separation of duties, and security architecture to name a few.

The *Security and Compliance Configuration Guide for VMware Cloud Foundation* outlines the steps to implement non-default configurations. Default configurations are confirmed and excluded from the configuration guide as part of the VMware Cloud Foundation post deployment steps. You must perform the procedures from the guide to ensure that the SDDC performance is not compromised.

The *VMware Cloud Foundation Audit Guide Appendix* supports the post-implementation process and audit process. It includes procedures to validate both default and non-default configurations. In the *VMware Cloud Foundation Audit Guide Appendix*, mappings between configurations and compliance controls provide a comprehensive inventory of configurations designated as default or non-default.

VMware Cloud Foundation Compliance Kit

Compliance kits apply to core products in VMware Cloud Foundation:

- VMware ESXi™
- VMware vCenter Server®
- VMware vSAN™
- VMware NSX™ Data Center
- VMware Cloud Foundation™ SDDC Manager

Default Access Controls Configured in VMware Cloud Foundation

Each product can support a range of settings that must be evaluated and if necessary, modified to meet security and compliance requirements.

Frequently requested access control settings are listed with the default values in VMware Cloud Foundation 5.2. Configurations with a value of 0 are deactivated.

The default settings may not be the recommended values based on desired the compliance standard. This is the default out-of-the-box state of access controls in VMware Cloud Foundation.

Table 255: Default Access Control Parameters in VMware Cloud Foundation

Product	Configuration ID	Configuration Description	Default Setting
NSX	VMW-NSXT-1416	Configure NSX-T Manager to terminate idle sessions after a certain period of time.	1800 seconds
NSX	VMW-NSXT-1417	Configure NSX-T Manager to block any login attempts after consecutive invalid login attempts for a certain period.	900 seconds
NSX	VMW-NSXT-1418	Configure NSX-T Manager to block further login attempts after a number of consecutive failed login attempts.	5 attempts
NSX	VMW-NSXT-1419	Configure NSX-T Manager locked accounts to automatically get unlocked after a period of time following the last failed login attempt.	900 seconds
NSX	VMW-NSXT-1421	Configure a minimum password length for NSX-T Manager accounts.	12 characters
ESXi	VMW-ESXI-00034	Set the maximum number of failed login attempts before an account is locked.	5 attempts
ESXi	VMW-ESXI-00038	Configure the inactivity timeout to automatically terminate idle shell sessions.	0 seconds (automatic termination is deactivated)
ESXi	VMW-ESXI-00109	Configure the password history to restrict the reuse of a certain number of previous passwords.	5
ESXi	VMW-ESXI-00165	Configure a time for automatic unlock of a locked user account.	900 seconds
ESXi	VMW-ESXI-00564	Configure the inactivity timeout to automatically terminate idle Host Client sessions.	900 seconds
ESXi	VMW-ESXI-00168	Configure the inactivity timeout to automatically terminate idle DCUI sessions.	600 seconds
vCenter Server	VMW-VC-00403	Configure the password history to restrict the reuse	5 passwords

Table continued on next page

Continued from previous page

Product	Configuration ID	Configuration Description	Default Setting
		of a certain number of previous passwords.	
vCenter Server	VMW-VC-00421	Configure vCenter Server to enforce a maximum password lifetime restriction.	90 days
vCenter Server	VMW-VC-00422	Configure the inactivity timeout to automatically terminate vSphere Client sessions.	120 minutes
vCenter Server	VMW-VC-00428	Configure vCenter Server to rotate the vpxuser auto-password periodically.	30 days
vCenter Server	VMW-VC-00427	Configure a minimum password length for the vpxuser account.	32 characters
vCenter Server	VMW-VC-00410	Configure the minimum number of characters for password length for any vCenter Server user.	8 characters
vCenter Server	VMW-VC-00408	Configure the minimum number of uppercase characters in the password for any vCenter Server user.	1 character
vCenter Server	VMW-VC-00413	Configure the minimum number of lowercase characters in the password for any vCenter Server user.	1 character
vCenter Server	VMW-VC-00433	Configure the minimum number of numeric characters in the password for any vCenter Server user.	1 character
vCenter Server	VMW-VC-00432	Configure the minimum number of special characters in the password for any vCenter Server user.	1 character
vCenter Server	VMW-VC-01271	Configure the maximum number of identical adjacent characters policy.	3 characters
vCenter Server	VMW-VC-00436	Limit the maximum number of failed login attempts for vCenter Server users.	5 attempts
vCenter Server	VMW-VC-00434	Configure the number of failed login attempts in a period of time before an account gets locked.	180 seconds
vCenter Server	VMW-VC-00435	Configure a timer for automatic account unlock	300 seconds

Table continued on next page

Continued from previous page

Product	Configuration ID	Configuration Description	Default Setting
		for accounts locked after failed login attempts.	
vCenter Server	VMW-VC-00096	Deactivate console connection sharing on the virtual machine.	1 (deactivated)

Security and Compliance Configuration Guide for VMware Cloud Foundation 5.2

Security and Compliance Configuration for VMware Cloud Foundation provides general guidance and step-by-step configuration for securing the management and workload domains in your VMware Cloud Foundation environment towards compliance with the NIST 800-53 standard. This guide is validated for the management workload domain and VI workload domains for VMware Cloud Foundation 5.2.

WARNING

This document is intended to provide general guidance for organizations that are considering VMware solutions to help them address compliance requirements. The information contained in this document is for educational and informational purposes only. This document is not intended to provide regulatory advice and is provided "AS IS". VMware makes no claims, promises or guarantees about the accuracy, completeness, or adequacy of the information contained herein. Organizations should engage appropriate legal, business, technical, and audit expertise within their specific organization for review of regulatory compliance requirements.

Intended Audience

Security and Compliance Configuration for VMware Cloud Foundation is intended for cloud architects, infrastructure administrators, and cloud administrators who are familiar with and want to use VMware software to secure and work towards compliance.

Required VMware Software

The *Security and Compliance Configuration for VMware Cloud Foundation* documentation is compliant and with certain product versions. See *VMware Cloud Foundation Release Notes* for more information about supported product versions.

Compliance Kit for VMware Cloud Foundation 4.5

The compliance kits is a solution that builds on top of VMware Cloud Foundation and leverages security fundamentals. The kit address the top ten most frequently requested compliance standards, regulations, and frameworks.

You can directly download the [VMware Cloud Foundation 5.2 Compliance Kit](#).

Update History

This *Security and Compliance Configuration for VMware Cloud Foundation* is updated with each release of the product or when necessary.

Revision	Description
23 JULY 2024	Initial release.

Software Requirements

To configure your VMware Cloud Foundation instance for compliance, you must download and license additional VMware and third-party software.

Security and Compliance Configuration for VMware Cloud Foundation uses scripts and commands based on VMware PowerCLI to reconfigure the SDDC. You must prepare a host with supported OS for running Microsoft PowerShell, set up Microsoft PowerShell, and install the latest version of VMware PowerCLI. The host must have connectivity to the ESXi management network in the management cluster.

Table 256: Additional Software Required for *Security and Compliance Configuration for VMware Cloud Foundation*

Product Group	Script/Tool	Description
VMware PowerCLI	Supported OS for VMware PowerCLI	Operating system that supports Microsoft PowerShell and VMware PowerCLI. For more information on supported operating systems, see VMware PowerCLI User's Guide .
VMware vSAN	Native Key Provider (NKP) or Key Management Server (KMS)	<p>If you are not using Native Key Provider (NKP) for encryption, Deploy and configure Key Management Server (KMS).</p> <p>Key Management Servers are developed and released by security and cloud vendors for encryption in virtualized environments. You use a Key Management Server to activate the encryption of vSAN storage. For a list of supported Key Management Server, see KMS list. Refer to the Key Management Server vendor documentation for setup and configuration instructions. Ensure that all encryption keys are available across regions to activate decryption in case of a region failover.</p>
VMware vSAN	Proxy server	vSAN uses an external proxy server to connect to the Internet to download the Hardware Compatibility List.
VMware NSX	SFTP server	Space for NSX Manager backups must be available on an SFTP server. The NSX Manager instances must have connection to the remote SFTP server.

Table 257: VMware Scripts and Tools Required for *Security and Compliance Configuration for VMware Cloud Foundation*

Product Group	Script/Tool	Download Location	Description
VMware vSphere	VMware PowerCLI	n/a	VMware PowerCLI contains modules of cmdlets based on Microsoft PowerShell for automating vSphere, VMware VMware NSX, and others. VMware PowerCLI provides a PowerShell interface to the VMware product APIs.

Securing ESXi Hosts

You perform procedures on the ESXi hosts in all your workload domains by using different interfaces, such as PowerCLI, ESXi Shell, and the vSphere Client.

Security Best Practices for Securing ESXi Hosts

You must follow multiple best practices at all times when you operate your ESXi hosts.

Table 258: Security Best Practices for Securing ESXi Hosts

Best Practice	Description
Add only system accounts to the ESXi exception users list. VMW-ESXI-00125	You can add users to the exception users list from the vSphere Client. These user accounts do not lose their permissions when the host enters lockdown mode. Only add service accounts such as backup agents. Do not add administrative users or user groups to exception users list. Adding unnecessary users to the exception list defeats the purpose of lockdown mode.
Install security patches and updates for ESXi hosts. VMW-ESXI-00129	You install all security patches and updates on the ESXi hosts as soon as the update bundles are available in SDDC Manager. Do not apply patches to ESXi manually or by using vSphere Update Manager or VMware vCenter Lifecycle Manager in a VMware Cloud Foundation environment unless directed to do so by support. If you patch the environment without using SDDC Manager, it can not only lead to a less-secure environment, but may cause issues with automated upgrades or actions in the future.
The ESXi host must protect the confidentiality and integrity of transmitted information by protecting ESXi management traffic. VMW-ESXI-00178	The vSphere management network provides access to the vSphere management interface on each component. Services running on the management interface provide an opportunity for an attacker to gain privileged access to the systems. Any remote attack most likely would begin with gaining entry to this network. The Management VMkernel port group must be on a dedicated VLAN. The Management VLAN must not be shared by any other function and must not be accessible to anything other than management-related functions such as vCenter.
The ESXi host must use approved certificates. VMW-ESXI-01113	The default self-signed, VMCA-issued host certificate must be replaced with a certificate from a trusted Certificate Authority (CA) when the host is accessed directly, such as during a virtual machine (VM) console connection.

Configure Multiple Security Settings on the ESXi Hosts by Using the ESXi Shell

You activate secure boot on all the ESXi hosts.

You perform the procedure from an ESXi Shell session connected to the ESXi host and on all ESXi hosts in the respective workload domain.

1. Log in to an ESXi host by using ESXi Shell as **root**.
2. VMW-ESXI-01108 Activate secure boot on the host.

```
/usr/lib/vmware/secureboot/bin/secureBoot.py -c
```

If the output indicates that secure boot cannot be activated, correct the discrepancies and try again. Once all discrepancies are resolved, the server ESXi is installed on can be updated to enable Secure Boot in the firmware.

To enable Secure Boot in the server's firmware, follow the instructions for the specific manufacturer.

3. Perform the procedure on the remaining hosts in the current and any other workload domains.

Configure Multiple Security Settings on the ESXi Hosts by Using PowerCLI

You perform the procedure on all ESXi hosts in all your workload domains to configure firewall settings, password policy, inactivity timeouts, failed login attempts, start up policies, join ESXi hosts to Active Directory domain, and remove ESX Admin group membership. Also, configure login banners for the Direct Console User Interface (DCUI) and SSH connections, deactivate warnings, configure persistent log location, remote logging, implement secure boot enforcement, enable TPM-based configuration encryption, enable audit logging, allocate storage record capacity, and activate bidirectional CHAP authentication by using PowerCLI commands.

To perform the procedure on the ESXi hosts for a workload domain, you connect to the vCenter Server for the respective workload domain. To run a task on all hosts for the domain, when you run commands, on the prompts to specify the object of a command, enter [A] Yes to all.

1. Log in to the vCenter Server for the workload domain you want to reconfigure by using a PowerCLI console.

Setting	Value
Command	Connect-VIServer -Server <i>management-domain-vcenter-server-fqdn</i> -Protocol https
User name	administrator@vsphere.local

2. VMW-ESXI-00022 Configure the password complexity policy for the ESXi host.

The requirement is a length of minimum 15 characters (maximum of 64 characters) from 4 character classes that include lowercase letters, uppercase letters, numbers, special characters. Password difference is also mandatory.

```
Get-VMHost | Get-AdvancedSetting -Name Security.PasswordQualityControl | Set-AdvancedSetting -Value "similar=deny retry=3 min=disabled,disabled,disabled,disabled,15 max=64"
```

3. VMW-ESXI-00028 Configure the ESXi hosts firewall to only allow traffic from the authorized management networks.

```
$esxiHosts = Get-VMHost

foreach($esxiHost in $esxiHosts){

$esxcli = Get-EsxCli -v2 -VMHost $esxiHost.Name

#This disables the allow all rule for the target service.The sshServer service is the target in this example.

$arguments = $esxcli.network.firewall.ruleset.set.CreateArgs()

$arguments.rulesetid = "sshServer"
```

```

$arguments.allowedall = $false

$esxcli.network.firewall.ruleset.set.Invoke($arguments)

#Next add the allowed IPs for the service. Note that executing the "vSphere Web
Client" service this way may disable access but may be done through vCenter or
through the console.

$arguments = $esxcli.network.firewall.ruleset.allowedip.add.CreateArgs()

$arguments.rulesetid = "sshServer"

$arguments.ipaddress = "Site-specific networks"

$esxcli.network.firewall.ruleset.allowedip.add.Invoke($arguments)

```

NOTE

This must be done for each user-configurable enabled service.

4. VMW-ESXI-00034 Set the maximum number of failed login attempts before an account is locked to 3.

```

Get-VMHost | Get-AdvancedSetting -Name Security.AccountLockFailures | Set-
AdvancedSetting -Value 3

```

5. VMW-ESXI-00038 Configure the inactivity timeout to automatically close idle shell sessions to 600 seconds.

```

Get-VMHost | Get-AdvancedSetting -Name UserVars.ESXiShellInteractiveTimeOut | Set-
AdvancedSetting -Value 600

```

6. VMW-ESXI-00039 Configure the timeout to automatically stop ESXi shell and SSH services to 600 seconds.

```

Get-VMHost | Get-AdvancedSetting -Name UserVars.ESXiShellTimeOut | Set-
AdvancedSetting -Value 600

```

7. VMW-ESXI-00114 To eliminate the need to create and maintain multiple local user accounts, join ESXi hosts to an Active Directory (AD) domain.

```

Get-VMHost | Get-VMHostAuthentication | Set-VMHostAuthentication -JoinDomain -Domain
"domain name" -User "username" -Password "password"

```

NOTE

If any local user accounts exist, apart from **root** and local service accounts, you can delete the local user accounts by going to the ESXi host UI **Manage > Security & Users > Users**.

8. VMW-ESXI-00122 Configure the login banner for the DCUI of the ESXi host.

```

Get-VMHost | Get-AdvancedSetting -Name Annotations.WelcomeMessage | Set-
AdvancedSetting -Value "Site-Specific banner text"

```

9. VMW-ESXI-00123 Configure the login banner for the SSH connections.

```

Get-VMHost | Get-AdvancedSetting -Name Config.Etc.issue | Set-AdvancedSetting -Value
"Site-Specific banner text"

```

10. VMW-ESXI-00136 Configure a persistent log location for all locally stored logs.

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.logDir | Set-AdvancedSetting
-Value "New Log Location"
```

NOTE

Specify the log location as [datastorename] path_to_file, where the path is relative to the root of the volume, backing the datastore. For example, the path [storage1] /systemlogs maps to the path /vmfs/volumes/storage1/systemlogs.

The new location should not include a subfolder as enabling audit logging will create a folder and will fail if a subfolder is specified.

11. VMW-ESXI-00137 For a host added to Active Directory, use an Active Directory group instead of the default **ESX Admins** group for the *esxAdminsGroup* property on the ESXi hosts.

```
Get-VMHost | Get-AdvancedSetting -Name
Config.HostAgent.plugins.hostsvc.esxAdminsGroup | Set-AdvancedSetting -Value site
specific AD_Group
```

NOTE

Changing the group name does not remove the permissions of the previous group.

12. VMW-ESXI-00164 Configure a remote log server for the ESXi hosts.

NOTE

Use the following format when adding the remote log server. You can enter multiple, comma-separated values.

```
udp://<IP/FQDN>:514
```

```
tcp://<IP/FQDN>:514
```

```
ssl://<IP/FQDN>:1514
```

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.logHost | Set-AdvancedSetting
-Value "<site-specific syslog server hostname>"
```

13. VMW-ESXI-00564 Configure idle session timeout for the ESXi host client to 600 seconds.

```
Get-VMHost | Get-AdvancedSetting -Name UserVars.HostClientSessionTimeout | Set-
AdvancedSetting -Value "600"
```

14. VMW-ESXI-01102 Activate bidirectional CHAP authentication for iSCSI traffic.

```
Get-VMHost | Get-VMHostHba | Where {$_.Type -eq "iscsi"} | Set-VMHostHba
-ChapType Required -ChapName chap_name -ChapPassword password -MutualChapEnabled
$true -MutualChapName mutual_chap_name -MutualChapPassword mutual_password
```

15. VMW-ESXI-01121 Activate strict x509 verification for SSL syslog endpoints.

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.certificate.strictX509Compliance
| Set-AdvancedSetting -Value "true"
```

16. VMW-ESXI-01122 Activate volatile key destruction on the host.

```
Get-VMHost | Get-AdvancedSetting -Name Mem.MemEagerZero | Set-AdvancedSetting -Value "1"
```

17. VMW-ESXI-01123 Configure the host with an appropriate maximum password age.

```
Get-VMHost | Get-AdvancedSetting -Name Security.PasswordMaxDays | Set-AdvancedSetting -Value "90"
```

18. VMW-ESXI-01124 Enable TPM-based configuration encryption.

- Ensure the TPM 2.0 chip is enabled in the BIOS and the ESX UI does not show any errors.
- This setting cannot be configured until the TPM is properly enabled in firmware.
- Configuration encryption uses the physical TPM at install or upgrade time. If the TPM is added or enabled later, you must reconfigure the ESXi host to use the newly available TPM. After you enable TPM configuration encryption is enabled, you cannot disable it.

```
$esxiHosts = Get-VMHost

foreach($esxiHost in $esxiHosts){

$esxcli = Get-EsxCli -v2 -VMHost $esxiHost.Name

$arguments = $esxcli.system.settings.encryption.set.CreateArgs()

$arguments.mode="TPM"

$esxcli.system.settings.encryption.set.Invoke($arguments)

}
```

You must evacuate the host and gracefully reboot for changes to take effect.

19. VMW-ESXI-01125 The ESXi host must implement Secure Boot enforcement.

```
$esxiHosts = Get-VMHost

foreach($esxiHost in $esxiHosts){

$esxcli = Get-EsxCli -v2 -VMHost $esxiHost.Name

$arguments = $esxcli.system.settings.encryption.set.CreateArgs()

$arguments.requiresecureboot =$true

$esxcli.system.settings.encryption.set.Invoke($arguments)

}
```

You must evacuate the host and gracefully reboot for changes to take effect.

20. VMW-ESXI-01126 Configure the startup policy for the CIM service on the host to "off".

```
Get-VMHost | Get-VMHostService | Where {$_.Label -eq "CIM Server"} | Set-VMHostService -Policy Off
```

21. VMW-ESXI-01128 Deactivate the startup policy for the SNMP service on the host.

```
Get-VMHost | Get-VMHostService | Where {$_.Label -eq "SNMP Server"} | Set-VMHostService -Policy Off
```

22. VMW-ESXI-01152 The ESXi host must disable virtual hardware management network interfaces.

```
Get-VMHost | Get-AdvancedSetting -Name Net.BMCNetworkEnable | Set-AdvancedSetting
-Value 0
```

23. VMW-ESXI-01141 The ESXi host must allocate audit record storage capacity to store audit records.

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.auditRecord.storageCapacity |
Set-AdvancedSetting -Value 100
```

24. VMW-ESXI-01142 ESXi host must enable audit logging.

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.auditRecord.storageEnable | Set-
AdvancedSetting -Value "true"
```

VMW-ESXI-00136 and VMW-ESXI-01141 must be configured and validated before enabling audit logging.

25. VMW-ESXI-01143 ESXi host must off-load audit records via syslog.

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.auditRecord.remoteEnable | Set-
AdvancedSetting -Value "true"
```

26. VMW-ESXI-01145 ESXi host must forward audit records containing information to establish what type of events occurred.

```
Get-VMHost | Get-AdvancedSetting -Name Syslog.global.logLevel | Set-AdvancedSetting
-Value "info"
```

27. VMW-ESXI-01150 The ESXi host must deny shell access for the dcui account.

```
$esxiHosts = Get-VMHost
foreach($esxiHost in $esxiHosts){
$esxcli = Get-EsxCli -v2 -VMHost $esxiHost.Name
$arguments = $esxcli.system.account.set.CreateArgs()
$arguments.id = "dcui"
$arguments.shellaccess = "false"
$esxcli.system.account.set.invoke($arguments)
}
```

28. VMW-ESXI-01153 ESXi host must enforce the exclusive running of executables from approved VIBs.

```
Get-VMHost | Get-AdvancedSetting -Name VMkernel.Boot.execInstalledOnly | Set-
AdvancedSetting -Value True
```

29. VMW-ESXI-01154 Configure ESXi host to use approved encryption to protect the confidentiality of network sessions.

```
$esxiHosts = Get-VMHost
foreach($esxiHost in $esxiHosts){
$esxcli = Get-EsxCli -v2 -VMHost $esxiHost.Name
$arguments = $esxcli.system.tls.server.set.CreateArgs()
$arguments.profile = "NIST_2024"
```

```
$esxcli.system.tls.server.set.invoke($arguments)
}
```

A reboot is required to complete the process of changing profiles.

Configure Multiple Security Settings on Unassigned ESXi Hosts by Using PowerCLI

You perform this procedure on all unassigned ESXi hosts in the SDDC inventory to configure non-native VLAN ID, Virtual Guest Tagging (VGT), and unreserved VLAN ID on all the port groups on the standard switch.

These controls apply only to unassigned hosts in VMware Cloud Foundation. An unassigned host is a host that is commissioned but not assigned to a workload domain. Once the host is added to a VMware Cloud Foundation workload domain, the standard switch on the host is removed and the host is added to a distributed switch.

The following configurations address ESXi standard switches only. Distributed switches are addressed in the [Securing vCenter Server](#) section (see [Securing vCenter Server](#)). If your environment does not have ESXi hosts with standard switches, you can skip this procedure.

1. Log in to the unassigned ESXi host you want to reconfigure by using a PowerCLI console and provide the credentials.

```
Connect-VIServer -Server host-fqdn -Protocol https
```

2. VMW-ESXI-01104 Do not configure the port groups on standard switches to VLAN 4095 unless Virtual Guest Tagging (VGT) is required.

```
Get-VirtualPortGroup -Name "portgroup name" | Set-VirtualPortGroup -VlanId "New VLAN#"
```

Activate Normal Lockdown Mode on the ESXi Hosts

You activate normal lockdown mode on the ESXi hosts.

1. In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	https://vcenter-server-fqdn/ui
User name	administrator@vsphere.local

2. VMW-ESXI-00031 Activate normal lockdown mode on a host.
 - a) In the **Hosts and clusters** inventory, select an ESXi host.
 - b) Click **Configure**.
 - c) Under **System**, select **Security profile**.
 - d) In the **Lockdown mode** panel, click **Edit**.
 - e) In the **Lockdown mode** dialog box, select the **Normal** or **Strict** radio button and click **OK**.

NOTE

In strict lockdown mode, the Direct Console User Interface (DCUI) service is stopped. If the connection to vCenter Server is lost and the vSphere Client is no longer available, the ESXi host becomes inaccessible.

3. Repeat the procedure for all ESXi hosts in all workload domains.

Securing vCenter Server

You perform procedures on the vCenter Server in all your workload domains using different interfaces: PowerCLI and vSphere Client.

Security Best Practices for Securing vCenter Server

You must follow multiple best practices at all times when you operate your vCenter Server instances.

Table 259: Security Best Practices for Securing vCenter Server

Best Practice	Description
Assign correct roles to vCenter Server users. VMW-VC-00415	Users and service accounts must be assigned only privileges they require. To reduce risk of confidentiality, availability, or integrity loss, the least privilege principle requires that these privileges must be assigned only if needed.
Use unique service accounts for applications that connect to vCenter Server. VMW-VC-00401	Create a service account for each application that connects to vCenter Server. Grant only the required permissions for the application to run.
vCenter Server must restrict access to cryptographic permissions. VMW-VC-01211	These permissions must be reserved for cryptographic administrators where VM and/or vSAN encryption is in use. Catastrophic data loss can result from a poorly administered cryptography. Only the Administrator and any site-specific cryptographic group must have the following permissions: <ul style="list-style-type: none"> • Cryptographic Operations privileges • Global.Diagnostics • Host.Inventory.Add host to cluster • Host.Inventory.Add standalone host • Host.Local operations.Manage user groups
The vCenter Server must use LDAPS when adding an SSO identity source. VMW-VC-01229	To protect the integrity of LDAP communications, secure LDAP (LDAPS) must be explicitly configured when adding an LDAP identity source in vSphere SSO. When configuring an identity source and supplying an SSL certificate, vCenter Server enforces secure LDAP.
The vCenter Server must implement Active Directory authentication VMW-VC-01228	The vCenter Server must ensure users are authenticated with an individual authenticator prior to using a group authenticator. Using Active Directory for authentication provides more robust account management capabilities.
Backup the vCenter Native Key Providers with a strong password. VMW-VC-01239	The vCenter Native Key Provider acts as a key provider for encryption based capabilities, such as encrypted virtual machines, without requiring an external KMS solution. When activating this feature, a backup PKCS#12 file is created. If no password is provided during the backup process, the backup file can be used maliciously and compromise the environment.
Restrict access to the cryptographic role. VMW-VC-01210	The built-in Administrator role has the permission to perform cryptographic operations, such as Key Management Server (KMS) functions and encrypting and

Table continued on next page

Continued from previous page

Best Practice	Description
	decrypting virtual machine disks. This role must be reserved for cryptographic administrators, where virtual machine or vSAN encryption is required. All other vSphere administrators, who do not require cryptographic operations, must be assigned the No cryptography administrator role.
<p>The vCenter Server Machine SSL certificate must be issued by an appropriate certificate authority.</p> <p>VMW-VC-01205</p>	<p>The default self-signed, VMCA-issued vCenter reverse proxy certificate must be replaced with an approved certificate. The use of an approved certificate on the vCenter reverse proxy and other services assures clients that the service they are connecting to is legitimate and trusted.</p>
<p>Ensure that port mirroring is used legitimately.</p> <p>VMW-VC-01248</p>	<p>The vSphere VDS can mirror traffic from one port to another, allowing observation of traffic. Ensure that port mirroring is used legitimately.</p>
<p>Install security patches and updates for vCenter Server.</p> <p>VMW-VC-01253</p>	<p>You install all security patches and updates on vCenter Server instances as soon as possible. An attacker can exploit known vulnerabilities when attempting to attain access or elevate privileges. Mitigate the risk of breaches by updating vCenter Server instances first and then updating ESXi hosts.</p>
<p>Configure Key Encryption Keys (KEKs) to be re-issued at regular intervals for the vSAN encrypted datastores.</p> <p>VMW-VC-01213</p>	<p>Interview the SA to determine whether a procedure exists to perform a shallow re-key of all vSAN encrypted datastores at regular, site-defined intervals. This interval must be defined by the SA and the ISSO. If vSAN encryption is not in use, this is not applicable.</p>
<p>At a minimum, vCenter must provide an immediate, real-time alert to the system administrator (SA) and information system security officer (ISSO) of all audit failure events requiring real-time alerts.</p> <p>VMW-VC-01254</p>	<p>Ensure that the Central Logging Server is configured to alert the SA and ISSO, at a minimum, on any AO-defined events. Otherwise, this is a finding. If there are no AO-defined events, this is not a finding.</p>
<p>Remove unnecessary virtual hardware devices from the VM.</p> <p>VMW-VC-01257</p>	<p>Ensure that no device is connected to a virtual machine if it is not required. For example, serial and parallel ports are rarely used for virtual machines in a datacenter environment, and CD/DVD drives are usually connected only temporarily during software installation. USB devices, sound cards, and other unnecessary hardware may be introduced with migrations from VMware Workstation, Fusion, or through other tools. Any enabled or connected device represents a potential attack channel, through the possibility of device drivers that contain vulnerabilities, by granting the ability to introduce software or exfiltrate data to or from a protected environment.</p> <p>Note: Removing the CD-ROM device may impact VMware Tools installation and maintenance.</p>
<p>vCenter is a version that has not reached End of General Support status.</p>	<p>Ensure that vCenter Server is of a version that has not reached End of General Support status.</p>

Table continued on next page

Continued from previous page

Best Practice	Description
VMW-VC-01256	
vCenter must separate authentication and authorization for administrators. VMW-VC-01261	Many organizations do both authentication and authorization using a centralized directory service such as Active Directory. Attackers who compromise an identity source can often add themselves to authorization groups, and simply log into systems they should not otherwise have access to. Additionally, reliance on central identity systems means that the administrators of those systems are potentially infrastructure administrators, too, as they can add themselves to infrastructure access groups at will. The use of local SSO groups for authorization helps prevent this avenue of attack by allowing the centralized identity source to still authenticate users but moving authorization into vCenter itself.
The vCenter Server must configure the firewall to only allow traffic from authorized networks. VMW-VC-01276	Ensures that all incoming and outgoing network traffic is blocked unless explicitly allowed, reducing the attack surface and helping to prevent unauthorized access to the system. Note that outgoing/egress traffic is not blocked, nor are related/established connections, so vCenter Server will still be able to communicate with systems where it initiates the connection. Perimeter firewalls should be used to curtail those types of connections.

Configure Security Settings for vCenter Server from the vSphere Client

You perform the procedure on all vCenter Server instances to configure password policies, lockout policies, alarms, proxy, login banners, LDAP, and other configurations.

1. In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	<code>https://management-domain-vcenter-server-fqdn/ui</code>
User name	<code>administrator@vsphere.local</code>

2. Configure the password policies.
 - a) From the **Home** menu of the vSphere Client, click **Administration**.
 - b) Under **Single Sign-On**, click **Configuration**.
 - c) On the **Local accounts** tab, under **Password policy**, click **Edit**.
 - d) In the **Edit password policies** dialog box, configure the settings and click **Save**.

Configuration ID	Setting	Value
VMW-VC-00421	Maximum lifetime	60
VMW-VC-00410	Minimum Length	15

Table continued on next page

Continued from previous page

Configuration ID	Setting	Value
VMW-VC-01269	Maximum Length	64

3. Configure the lockout policies.

- On the **Local accounts** tab, under **Lockout policy**, click **Edit**.
- In the **Edit lockout policies** dialog box, configure the settings and click **Save**.

Configuration ID	Setting	Value
VMW-VC-00436	Maximum number of failed login attempts	3
VMW-VC-00434	Time interval between failures	900 seconds
VMW-VC-00435	Unlock time	0 seconds

4. VMW-VC-01219 Configure an alert for the appropriate personnel about SSO account actions

- In the **Hosts and clusters** inventory, select the vCenter Server that manages the ESXi host you configure.
- Click the **Configure** tab, select **Alarm definitions** under **Security**.
- Click **Add**.
The **New alarm definition** wizard opens.
- On the **Name and targets** page, enter the settings and click **Next**.

Setting	Value
Alarm name	SSO account actions - com.vmware.sso.PrincipalManagement
Target type	vCenter Server

- On the **Alarm rule 1** page, under **If**, enter `com.vmware.sso.PrincipalManagement` as a trigger and press Enter.
- Configure the remaining settings for the alarm, click **Next**, and follow the prompts to finish the wizard.

Setting	Value
Trigger the alarm and	Show as warning
Send email notifications	Off
Send SNMP traps	On
Run script	Off

5. VMW-VC-01209 Configure a login message.

- From the **Home** menu of the vSphere Client, click **Administration**.
- Navigate to **Single sign-on > Configuration**.
- Click the **Login message** tab and click **Edit**.
- Activate the **Show login message** toggle.
- In the **Login message** text box, enter the login message.
- Activate the **Consent checkbox** toggle.
- In the **Details of login message** text box, enter the site-specific banner text and click **Save**.

6. VMW-VC-01212 Configure Mutual CHAP for vSAN iSCSI targets.
 - a) In the **Hosts and Clusters** inventory, select the vSAN-enabled cluster.
 - b) Click the **Configure** tab and under **vSAN**, click **Services**.
 - c) In the **vSAN iSCSI target service** tile, click **Enable**.
 - d) Activate the service from the toggle switch.
 - e) From the **Authentication** drop-down menu, select **Mutual CHAP**
 - f) Configure the incoming and outgoing users and secrets appropriately and click **Apply**.
7. Set SDDC deployment details on the vCenter Server instances.
 - a) In the **Global inventory lists** inventory, click **vCenter Servers**.
 - b) Click the vCenter Server object and click the **Configure** tab in the central pane.
 - c) Under **Settings**, click **Advanced settings** and click **Edit settings**.
 - d) In the **Edit advanced vCenter Server settings** dialog box, enter the settings and click **Add**.

Setting	Value
Name	config.SDDC.Deployed.ComplianceKit
Value	VCF-NIST-800-53

8. VMW-VC-00422 vCenter Server must terminate vSphere Client sessions after 10 minutes of inactivity.
 - a) From the **Home** menu of the vSphere Client, click **Administration**.
 - b) Under **Deployment**, click **Client configuration**.
 - c) Click **Edit**, for **Session timeout** , enter 10 minutes, and click **Save**.
9. VMW-VC-01216 vCenter must limit membership to the SystemConfiguration.BashShellAdministrators SSO group.
 - a) From the **Home** menu of the vSphere Client, click **Administration**.
 - b) Under **Single sign-on**, click **Users and Groups** and **Groups**.
 - c) Click **> next page arrow** until **SystemConfiguration.BashShellAdministrators** appears.
 - d) Click `SystemConfiguration.BashShellAdministrators` and click three vertical dots next to the name of each unauthorized account and click **Remove Member** and click **Remove**.

NOTE

By default the Administrator and a unique service account similar to "vmware-applmgmtservice-714684a4-342f-4eff-a232-cdc21def00c2" will be in the group and should not be removed.

10. VMW-VC-01217 vCenter must limit membership to the TrustedAdmins SSO group.
 - a) From the **Home** menu of the vSphere Client, click **Administration**.
 - b) Under **Single sign-on**, click **Users and Groups** and **Groups**.
 - c) Click **> next page arrow** until **TrustedAdmins** appears.
 - d) Click `TrustedAdmins` and click three vertical dots next to the name of each unauthorized account and click **Remove Member** and click **Remove**.

NOTE

These accounts act as root on the Photon operating system and have the ability to severely damage vCenter, inadvertently or otherwise.

11. VMW-VC-01274 The vCenter Server must disable accounts used for Integrated Windows Authentication (IWA).
 - a) From the **Home** menu of the vSphere Client, click **Administration**.
 - b) Under **Single sign-on**, click **Users and Groups** and **Users**.
 - c) Change the domain to **vSphere.local** and

- d) Select **K/M** and **krbtgt/VSPHERE.LOCAL** accounts and click **More** and select **Disable** and click **OK**.
 - e) Repeat Step d with **krbtgt/VSPHERE.LOCAL** account
12. VMW-VC-01267 vCenter must require authentication for published content libraries.
- a) From the **Home** menu of the vSphere Client, click **Content Libraries**.
 - b) Click on the target content library and click **Edit Settings** under **Actions**.
 - c) Click the checkbox to **Enable user authentication for access to this content library**, and enter and confirm **password** and click **OK**.

NOTE

Any subscribed content libraries will need to be updated to enable authentication and provide the password.

13. VMW-VC-01268 vCenter must enable the OVF security policy for content libraries.
- a) From the **Home** menu of the vSphere Client, click **Content Libraries**.
 - b) Click on the target content library and click **Edit Settings** under **Actions**.
 - c) Click the checkbox to **Apply Security Policy**, and click **OK**.

NOTE

If you disable the security policy of a content library, you cannot reuse the existing OVF items.

Configure Security Settings for vCenter Server by Using PowerCLI

To configure host password length, native VLAN, reserved VLAN, and VGT, you perform the procedure on all vCenter Servers instances.

1. Log in to vCenter Server by using a PowerCLI console.

Setting	Value
Command	Connect-VIServer -Server <i>management-domain-vcenter-server-fqdn</i> -Protocol https
User name	administrator@vsphere.local

2. VMW-VC-01201 Configure all port groups to a value different from the value of the native VLAN.

```
Get-VDPortgroup "portgroup name" | Set-VDVlanConfiguration -VlanId "New VLAN#"
```
3. VMW-VC-01202 Configure all port groups to VLAN values not reserved by upstream physical switches

```
Get-VDPortgroup "portgroup name" | Set-VDVlanConfiguration -VlanId "New VLAN#"
```
4. VMW-VC-01227 Do not configure VLAN trunking in vCenter Server unless Virtual Guest Tagging (VGT) is required and authorized.
 - a) If you use VLAN ranges, enter VLAN ranges with a comma separated value to configure trunking.

```
Get-VDPortgroup "Portgroup Name" | Set-VDVlanConfiguration -VlanTrunkRange "<VLAN Range(s) comma separated>"
```
 - b) If you use a single VLAN, enter a single VLAN ID to configure trunking.

```
Get-VDPortgroup "Portgroup Name" | Set-VDVlanConfiguration -VlanId "<New VLAN#>"
```

5. VMW-VC-01247 **Services that may be unnecessary should be disabled such as CDP or LLDP network discovery protocols.**

```
Get-VDSwitch -Name "DSwitch" | Set-VDSwitch -LinkDiscoveryProtocolOperation  
"Disabled"
```

6. VMW-VC-01265 **vCenter must reset port configuration when virtual machines are disconnected.**

```
$pgs = Get-VDPortgroup | Get-View  
ForEach($pg in $pgs){  
$spec = New-Object VMware.Vim.DVPortgroupConfigSpec  
$spec.configversion = $pg.Config.ConfigVersion  
$spec.Policy = New-Object VMware.Vim.VMwareDVSPortgroupPolicy  
$spec.Policy.PortConfigResetAtDisconnect = $True  
$pg.ReconfigureDVPortgroup_Task($spec)  
}
```

7. VMW-VC-01266 **vCenter must not override port group settings at the port level on distributed switches, except for block ports.**

```
$pgs = Get-VDPortgroup | Get-View  
ForEach($pg in $pgs){  
$spec = New-Object VMware.Vim.DVPortgroupConfigSpec  
$spec.configversion = $pg.Config.ConfigVersion  
$spec.Policy = New-Object VMware.Vim.VMwareDVSPortgroupPolicy  
$spec.Policy.VlanOverrideAllowed = $False  
$spec.Policy.UplinkTeamingOverrideAllowed = $False  
$spec.Policy.SecurityPolicyOverrideAllowed = $False  
$spec.Policy.IpfixOverrideAllowed = $False  
$spec.Policy.BlockOverrideAllowed = $True  
$spec.Policy.ShapingOverrideAllowed = $False  
$spec.Policy.VendorConfigOverrideAllowed = $False  
$spec.Policy.TrafficFilterOverrideAllowed = $False  
$pg.ReconfigureDVPortgroup_Task($spec)  
}
```

8. VMW-VC-01275 **Configure the vCenter Server login banner text for access via SSH.**

```
Get-AdvancedSetting -Entity $VC -Name etc.issue | Set-AdvancedSetting -Value  
"Authorized login banner"
```

Configure Security Settings on the vCenter Server Appliance

You configure a syslog server and configure backups for vCenter Server from the vCenter Server Appliance Management Interface.

1. In a Web browser, log in to the vCenter Server Management Interface.

Setting	Value
URL	https://vcenter-server-fqdn:5480
User name	administrator@vsphere.local

2. VMW-VC-01218 Configure the appliance to send logs to a central log server.
 - a) In the left pane, click **Syslog**.
 - b) Click **Configure**, configure the address and port of a site-specific syslog aggregator or SIEM with the appropriate protocol, and click **Save**.

NOTE

UDP is discouraged due to it's stateless and unencrypted nature. TLS is recommended.

3. VMW-VC-01220 The vCenter Server configuration must be backed up on a regular basis.
 - a) In the left pane, click **Backup** and click **Configure** or **Edit** for an existing configuration.
 - b) Enter site-specific information for the backup job.
 - c) Ensure that the schedule is set to **Daily** and click **Create**.
4. In a Web browser, log in to the vCenter Server Management Interface.

Setting	Value
URL	https://vcenter-server-fqdn:5480
User name	root

5. VMW-VC-01255 Ensure password expiration for the root user is correct.
 - a) In the left pane, click **Administration** and click **Edit** under Password Expiration Settings.
 - b) Set **Password Validity (days)** to 90 and **Email for expiration warning** to your own email address and click **SAVE**.

NOTE

Configure SMTP on vCenter Server to receive the notification of expiration warning.

Securing SDDC Manager

You perform the procedures on SDDC Manager instances in your environment.

Security Best Practices for Securing SDDC Manager

You must follow multiple best practices at all times when you operate your SDDC Manager instances.

Table 260: Security Best Practices for Securing SDDC Manager

Best Practice	Description
Verify SDDC Manager backup VMW-SDDC-01600	<p>You must back up SDDC Manager regularly to avoid downtime and data loss in case of a system failure. You can back up and restore SDDC Manager with an image-based or a file-based solution. File-based backup is recommended for customers who are comfortable with configuring backups by using APIs, and are not using composable servers or stretched clusters.</p> <p>For image-based backups of SDDC Manager, use a solution compatible with VMware vSphere Storage APIs - Data Protection.</p> <p>For file-based backups, configure an external SFTP server as a target backup location and configure a backup schedule.</p>
The SDDC Manager must sync internal clocks with an authoritative time source. VMW-SDDC-01601	<p>Determining the correct time a particular application event occurred on a system is critical when conducting forensic analysis and investigating system events.</p> <p>Synchronization of system clocks is needed in order to correctly correlate the timing of events that occur across multiple systems. To meet this requirement, the organization will define an authoritative time source and have each system compare its internal clock at least every 24 hours. From the SDDC Manager UI, navigate to Administration >> Network Settings >> NTP Configuration to configure NTP server.</p>
Install security patches and updates for SDDC Manager VMW-SDDC-01602	Install all security patches and updates. To apply patches and updates to SDDC Manager, follow the guidance in the <i>VMware Cloud Foundation Lifecycle Management</i> document.
Use SSL certificates issued by a trusted certificate authority for SDDC Manager VMW-SDDC-01603	The use of a trusted certificate on the SDDC Manager appliance assures clients that the service they are connecting to is legitimate and trusted. To update the SDDC Manager certificate, refer the following URL: Install Certificates with External or Third-Party Certificate Authorities .
Do not expose SDDC Manager directly to the internet VMW-SDDC-01604	Allowing external access to the SDDC Manager appliance can expose the server to denial of service attacks or other penetration attempts. System Administrator (SA) should work with the network or boundary team to ensure proper firewall rules are configured or other mechanisms are in place to protect the SDDC Manager appliance.
Assign least privileges to users and service accounts in SDDC Manager VMW-SDDC-01605	<p>Users and groups must be assigned only privileges they require. To reduce risk of confidentiality, availability, or integrity loss, least privilege requires that these privileges must be assigned only if needed.</p> <p>From the SDDC Manager UI, under Administration > Single Sign On > Users and groups, review the users and</p>

Table continued on next page

Continued from previous page

Best Practice	Description
	groups assigned a role in SDDC Manager and verify that an appropriate role is assigned.
Dedicate an account for downloading updates and patches in SDDC Manager VMW-SDDC-01607	When access is allowed to download updates online, using a dedicated My VMware account ensures consistent access to updates and security patches in the event of system administrator turnover or account access issues. To configure a dedicated account that is not associated with a particular system administrator, from the SDDC Manager UI, go to Administration > Depot Settings .
Deploy SDDC Manager with FIPS security mode activated VMW-SDDC-01608	FIPS mode must be activated during bring-up and cannot be activated post bring-up. Refer to the VCF deployment guide for details on activating FIPS mode on SDDC Manager. CAUTION This option is only available for new VMware Cloud Foundation installations and the setting you apply during bring-up are used for future upgrades. You cannot change the FIPS security mode setting after bring-up.

Configure Security Settings for SDDC Manager by Using the SDDC Manager UI

To configure automatic password rotation, you perform the procedure in the SDDC Manager UI .

If you change the vCenter Server password length using the vSphere Client or the ESXi password length using the VMware Host Client, rotating the password for those components by using SDDC Manager generates a password that complies with the password length that you specified.

Automatic password rotation is currently not supported for ESXi.

SDDC Manager has default password policy settings for automatic password rotation.

Table 261: Default Password Settings for Automatic Password Rotation by SDDC Manager

Setting	Value
Minimum length	20 characters
Minimum uppercase characters	1
Minimum numeric characters	1
Minimum special characters	1
Maximum consecutive identical characters	2

1. In a Web browser, log in to the SDDC Manager using the SDDC Manager UI.

Setting	Value
URL	https://sddc_manager-fqdn/ui

Table continued on next page

Continued from previous page

Setting	Value
User name	administrator@vsphere.local

2. VMW-SDDC-01609 Schedule automatic password rotation for vCenter Server, Platform Services Controller (PSC), VMware NSX, and, backup.
 - a) In the left pane, navigate to **Security > Password management**.
 - b) Select a component (such as vCenter).
 - c) Select the username(s), click **Schedule rotation**, and select a rotation schedule (30, 60, or 90 days).
 - d) Click "Yes" to confirm.

Securing Management Virtual Machines

You connect to the management domain vCenter Server and use a script to perform multiple configurations on the management virtual machines that belong to the management domain. vSphere Cluster Services (vCLS) nodes are not in scope of this procedure as they are service VMs.

To harden the management VMs, you must power off the VMs one by one and run the script. To harden the vCenter Server VM, follow the instructions below:

1. Disable the lockdown mode on the ESXi host that hosts vCenter Server VM.
2. PowerOff the vCenter Server VM.
3. Run the below script by connecting to ESXi Host using `Connect-VIServer -Server <ESXi host FQDN which hosts vCenter Server VM> cmdlet`.
4. Login to ESXi host client that hosts the vCenter Server VM.
5. Power on the vCenter Server VM.
6. Enable the lockdown mode on the ESXi host.

If ESXi is version 7.0 U3i or above, you can run the script without powering off the management VMs. You must shut down the guest OS and power on (cold boot) the VMs for the advanced settings to take effect. Do not reboot the VMs. To prevent service interruption, cold boot must be performed one virtual machine at a time. Cold boot of vCenter Server and SDDC Manager requires a maintenance window.

Perform cold boot in the following order:

1. NSX Edge nodes
2. NSX Manager nodes
3. vCenter Server
4. SDDC Manager

Configuration ID	Description
VMW-VC-00096	Limit console connection sharing

1. Log in to the management domain vCenter Server by using a PowerCLI console.

Setting	Value
Command	<code>Connect-VIServer -Server <i>management-domain-vcenter-server-fqdn</i> -Protocol https</code>

Table continued on next page

Continued from previous page

Setting	Value
User name	administrator@vsphere.local

2. Configure advanced settings on all management virtual machines by running the script.

You must enter the name of the VM that you are reconfiguring in the first line of the script. For example, `$VMs = ("sddc-manager")`. If ESXi is version 7.0 U3i, you can enter a comma separated list of VMs.

```
$VMs = (management-domain-VM-name)

Foreach ($vm in $VMs){
    $advancedSetting = "RemoteDisplay.maxConnections"

    $setting = Get-VM $vm | Get-AdvancedSetting -Name $advancedSetting | Select-Object -Property Name, Value

    if(!$setting.Name){
        Get-VM $vm | New-AdvancedSetting -Name $advancedSetting -Value 1 -Confirm:$false
    }

    elseif($setting.Value -ne 1){
        Get-VM $vm | Get-AdvancedSetting -Name $advancedSetting | Set-AdvancedSetting -Value 1 -Confirm:$false
    }
}
```

Securing vSAN

You perform procedures on the vCenter Server instance by using the vSphere Client.

Security Best Practices for Securing vSAN

You must follow multiple best practices at all times when you operate your vSAN storage.

Table 262: Security Best Practice for Securing vSAN

Best Practice	Description
vSAN must reserve space to complete internal maintenance operations. VMW-vSAN-00186	vSAN Operations Reserve capacity setting helps ensure that vSAN always has sufficient free space to maintain the availability and reliability of the vSAN datastore and prevent potential data loss or service disruptions due to insufficient capacity during operations like policy changes.

Table continued on next page

Continued from previous page

Best Practice	Description
	This configuration parameter can be altered while the cluster is operational.
NFS file shares on vSAN File Services must be configured to restrict access. VMW-vSAN-00185	When configuring an NFS file share the "Customize net access" option should be selected with a restrictive set of permissions configured.
SMB file shares on vSAN File Services must accept only encrypted SMB authentication communications. VMW-vSAN-00187	When configuring an SMB file share the Protocol Encryption option must be enabled.

Configure a Proxy Server for vSAN from the vSphere Client

You perform the procedure on the respective vCenter Server to configure a proxy server for the download of the public Hardware Compatibility List.

1. In a Web browser, log in to your vCenter Server by using the vSphere Client.

Setting	Value
URL	<code>https://management-domain-vcenter-server-fqdn/ui</code>
User name	<code>administrator@vsphere.local</code>

2. VMW-vSAN-00207 Configure a proxy for the download of the public Hardware Compatibility List.
 - a) In the **Hosts and Clusters** inventory, select the vCenter Server object.
 - b) Click the **Configure** tab and under **vSAN**, click **Internet connectivity**.
 - c) On the **Internet connectivity** page, click **Edit**.
 - d) Select the **Configure the proxy server if your system uses one** check box.
 - e) Enter the proxy server details and click **Apply**.

Securing VMware NSX

You perform the procedures on different components of NSX.

Security Best Practices for Securing VMware NSX

You must follow multiple best practices at all times when you operate your NSX environment.

Table 263: NSX

Best Practice and Configuration ID	Description
Use roles and privileges in NSX Manager to limit user privileges. VMW-NSX-01410	Users and service accounts must be assigned the required privileges only. You can create a new role with reduced permissions. Navigate to System > Settings > User management > Rol

Table continued on next page

Continued from previous page

Best Practice and Configuration ID	Description
	<p>es. Click Add role, provide a name, the required permissions, and click Save.</p> <p>You can reduce permissions to an existing role. Navigate to System › Settings › User Management › User role assignment. Click the vertical ellipsis next to the target user or group, select Edit, remove the existing role, select the new role, and click Save.</p>
Integrate VMware Identity Manager (vIDM) or OpenID Connect (which supports multi factor authentication) with NSX VMW-NSX-01415	Use vIDM or OpenID Connect to meet requirements for authentication, authorization, and access control.
NSX Manager must obtain its public key certificates from an approved certificate authority. VMW-NSX-01466	For user certificates, each organization obtains certificates from an approved, shared service provider, as required by OMB policy. For federal agencies operating a legacy public key infrastructure cross-certified with the Federal Bridge Certification Authority at medium assurance or higher, this Certification Authority will suffice.

Configure Security Settings for VMware NSX by Using the User Interfaces

You perform the procedure in NSX to configure logging servers, configure logging for distributed and gateway firewall rules, and configure port binding for the spoofguard profile. Configure the settings for all NSX instances in your VMware Cloud Foundation environment.

1. In a Web browser, log in to the NSX Manager cluster as an administrator by using the user interface.
2. VMW-NSX-01468 You configure NSX Manager to perform backups on an organizational defined schedule.
 - a) On the main navigation bar, click **System**.
 - b) In the left pane, navigate to **Lifecycle management › Backup and restore**.
 - c) Next to **SFTP server**, click **Edit**.
 - d) In the **Backup configuration** dialog box, enter the required details and click **Save**.
 - e) Next to **Schedule**, click **Edit**.
 - f) In the **Schedule recurring backup** dialog box, click **Recurring backup toggle** and configure an interval between backups.
 - g) To perform backups on detection of configuration changes, activate **Detect NSX configuration change**, specify an interval for detecting changes, and click **Save**.
3. VMW-NSX-01500 The NSX Manager must disable unused local accounts.
 - a) On the main navigation bar, click **System**.
 - b) In the left pane, navigate to **Settings › User management**.
 - c) Click **Local users** and click vertical ellipsis next to the user to modify and click **Deactivate User**.
4. VMW-NSX-01524 NSX Manager must display the Standard Mandatory Notice and Consent Banner before granting access.
 - a) On the main navigation bar, click **System**.
 - b) In the left pane, navigate to **Settings › General Settings**.
 - c) Click **User Interface** and click **Edit** next to **Login Consent Settings**.
 - d) Toggle **Login Consent** to On and **Require Explicit User Consent** to Yes.
 - e) Input **Consent Message Title** with Standard mandatory notice and consent banner and **Consent Message Description** and click **Save**.

Configure Security Settings for NSX by Using CLI Commands

You configure NSX Manager to back up audit records to a logging server. Also, you configure NSX Edge nodes to back up audit records to a central audit server.

1. VMW-NSX-01401 Synchronize internal information system clocks using redundant authoritative time sources.

- a) Open the VM console of the NSX Manager appliance in vCenter Server and log in with credentials authorized for administration.
- b) Run the following commands:

```
#remove any unknown or nonauthoritative NTP servers

del ntp-server <server-ip or server-name>

#configure ntp server

set ntp-server <server-ip or server-name>
```

2. VMW-NSX-01414 Configure NSX Manager to send logs to a central log server.

You can configure the logging server with one of the following protocols: TCP, LI-TLS, or TLS. If you use the protocols TLS or LI-TLS to configure a secure connection to a log server, the server and client certificates must be stored in the `/image/vmware/nsx/file-store/` folder on each NSX Manager appliance.

- a) Open the VM console of the NSX Manager appliance in vCenter Server and log in with credentials authorized for administration.
- b) If you want to configure a TCP or UDP syslog server, run `set logging-server <server-ip_or_server-name> proto <tcp or udp> level info` and press Enter.
- c) If you want to configure a TLS syslog server, run `set logging-server <server-ip_or_server-name> proto tls level info serverca ca.pem clientca ca.pem certificate cert.pem key key.pem` and press Enter.
- d) If you want to configure an LI-TLS server, run `set logging-server <server-ip_or_server-name> proto li-tls level info serverca root-ca.crt` and press Enter.

3. VMW-NSX-01421 Enforce a minimum of 15 characters for password length on the NSX Manager nodes.

- a) Open the VM console of an NSX Manager appliance in vCenter Server and log in with credentials authorized for administration.
- b) Run the command and press Enter.

```
set password-complexity minimum-password-length 15
```

4. VMW-NSX-01530 NSX Manager must require that when a password is changed, the characters are changed in at least eight of the positions within the password.

- a) Open the VM console of an NSX Manager appliance in vCenter Server and log in with credentials authorized for administration.
- b) Run the command and press Enter.

```
set password-complexity max-repeats 8
```

5. Configure login sessions settings for the NSX Manager.

- a) Open the VM console of the NSX Manager appliance in vCenter Server and log in with credentials authorized for administration.
- b) VMW-NSX-01416 Configure session lock after a 10-minute period of inactivity.

```
Set service http session-timeout 600
```

- c) VMW-NSX-01418 Prevent an account from further log in attempts by using the UI or API after three consecutive failed log in attempts.

```
Set auth-policy api max-auth-failures 3
```

- d) VMW-NSX-01498 Prevent an account from further log in attempts by using CLI after three consecutive failed log in attempts.

```
set auth-policy cli max-auth-failures 3
```

Configure Security Settings for VMware NSX by Using NSX API

You configure TLS 1.2 protocol and disable TLS 1.1 for NSX Manager.

1. VMW-NSX-01501 Configure an NSX Manager node to only use the TLS 1.2 protocol.

The change applies to all nodes in the cluster. The API service on each node restarts after the update. A delay of up to a minute between the time this API call completes and when the new configuration applies is possible.

- a) Run the GET command and save the output.

```
GET https://<nsx-mgr>/api/v1/cluster/api-service
```

- b) In the saved output, edit the `protocol_versions` line to disable TLS 1.1.

```
        "protocol_versions": [ { "name": "TLSv1.1", "enabled": false },
{ "name": "TLSv1.2", "enabled": true } ]
```

- c) Run the API call using curl or another REST API client with the edited initial output.

```
PUT https://<nsx-mgr>/api/v1/cluster/api-service
```

Optional Security Configurations for VMware NSX

Application Virtual Networks (AVN)s, which include the NSX Edge Cluster and NSX network segments, are no longer deployed and configured during bring-up. Instead they are implemented as a Day-N operations in SDDC Manager, providing greater flexibility. These configurations should be reevaluated if you plan to deploy NSX edges in your environment. Similarly, Distributed Firewall (DFW) and Gateway Firewall configurations are applicable only if you purchase NSX Firewall Add-On. These configurations are included as optional and site-specific only.

Configure Security Settings for NSX Edge Nodes by Using the User Interface

You perform the procedure in NSX to configure traffic logging for Gateway Firewall rules, publish any firewall policy/rule changes, deny traffic by default, flood protection profile, ingress filters, restrict traffic and disable Internet Control Message Protocol (ICMP) unreachable notifications, mask replies, redirects on the external interfaces. Configure the settings for all NSX edge instances in your VMware Cloud Foundation environment.

1. In a Web browser, log in to the NSX Manager cluster as an administrator by using the user interface.
2. VMW-NSX-01429, VMW-NSX-01514 Configure the NSX Gateway Firewall on the tier-0 and tier-1 gateways to generate traffic log entries.

NOTE

If the tier-0 gateway is deployed in an active/active high availability mode and no stateless rules exist, this configuration is not applicable.

- a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Policy Management > Gateway Firewall/**
 - c) Click the **Gateway specific rules** tab.
 - d) From the **Gateway** drop-down menu, select the respective gateway.
 - e) For each tier-0 gateway and for each rule with logging disabled, click the gear icon, activate the **Logging** toggle, and click **Apply**.
 - f) On the **Gateway Firewall** page, click **Publish**.
 - g) Repeat the procedure for each tier-1 gateway and for each rule with deactivated logging.
3. VMW-NSX-01431, VMW-NSX-01432 Configure the NSX Gateway Firewall on the tier-0 and tier-1 gateways to deny network traffic by default and allow network traffic by exception.
- a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Policy Management > Gateway Firewall**.
 - c) Click the **Gateway specific rules** tab.
 - d) From the **Gateway** drop-down menu, select the respective gateway.
 - e) Expand the default policy, and from the **Actions** drop-down menu, select **Reject** or **Drop**.
 - f) On the **Gateway Firewall** page, click **Publish**.
 - g) Repeat the procedure for each tier-1 gateway.
4. VMW-NSX-01437 Configure the multicast NSX tier-0 gateway to deactivate Protocol Independent Multicast (PIM) on all interfaces that are not required to support multicast routing.
- a) Navigate to **Networking > Connectivity > Tier-0 Gateways** and expand the target Tier-0 gateway.
 - b) Expand **Interfaces and GRE Tunnels**, click on the number of **External and Service interfaces** present to open the interfaces dialog, and then select "Edit" on the target interface.
 - c) Expand "Multicast", change PIM to "Deactivated", and then click "Save".
5. VMW-NSX-01438 Remove inactive interfaces on an NSX Tier-0 gateway.
- a) Navigate to **Networking > Connectivity > Tier-0 Gateways** and expand the target Tier-0 gateway.
 - b) Expand **Interfaces and GRE Tunnels**, click on the number of "External and Service interfaces" present to open the interfaces dialog, and then select "Edit" on the target interface.
 - c) Select "Delete" on the unneeded interface, and then click "Delete" again to confirm.
6. VMW-NSX-01442 Disconnect inactive linked segments for NSX Tier-1 gateways.
- a) Navigate to **Networking > Connectivity > Segments** and edit the target segment.
 - b) Under Connected Gateway, change to "None" and click "Save".

NOTE

The stale linked segment can also be deleted if there are no active workloads attached to it.

- c) Navigate to **Networking > Connectivity > Tier-1 Gateways** and edit the target Tier-1 Gateway.
 - d) Expand Service Interfaces and click on the number to view the Service Interfaces.
 - e) On the stale service interface, select **Delete** and click **Delete** again to confirm.
7. VMW-NSX-01453, VMW-NSX-01515 Configure flood protection profiles on the NSX Gateway Firewall for the tier-0 and tier-1 gateways to protect against Denial of Service (DDoS) attacks.

NOTE

If the tier-0 gateway is deployed in an active/active high availability mode and no stateless rules exist, this configuration is not applicable.

- a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Settings > General Settings**.
 - c) Click the **Firewall > Flood Protection** under **General Security Settings** tab.
 - d) From the **Add profile** drop-down menu, select **Add Edge Gateway profile**.
 - e) Enter a name and specify appropriate values for the following: **TCP half open connection limit**, **UDP active flow limit**, **ICMP active flow limit**, and **Other active connection limit**.
 - f) Configure the **Applied to** field to contain the tier-0 gateways, and then click **Save**.
 - g) Repeat this step for the tier-1 gateway and set **Applied to** to contain the tier-1 gateways.
8. VMW-NSX-01460 To protect against route table flooding and prefix de-aggregation attacks, configure the NSX tier-0 gateway to use maximum prefixes.
- a) On the main navigation bar, click **Networking**.
 - b) In the left pane, navigate to **Connectivity > Tier-0 gateways**.
 - c) Expand the NSX tier-0 gateway.
 - d) Expand the **BGP** section and click **BGP neighbors**.
 - e) In the **Set BGP neighbors** dialog box, click the vertical ellipsis and click **Edit** for the first neighbor.
 - f) Click the number in the **Route filter** column.
 - g) To configure the maximum routes value, specific to your environment, in the **Set route filter** dialog box, click the vertical ellipsis menu and click **Edit**.
 - h) Repeat the step to configure all neighbors.
9. VMW-NSX-01494, VMW-NSX-01495, VMW-NSX-01496 Configure the NSX tier-0 gateway to have Internet Control Message Protocol (ICMP) unreachable notifications, mask replies, and disable redirects on all external interfaces.

NOTE

If the tier-0 gateway is deployed in an active/active high availability mode and no stateless rules exist, this configuration is not applicable.

NSX does not come with a pre-configured service for ICMP mask replies. You may need to create this service.

- a) On the main navigation bar, click **Security**.
- b) In the left pane, navigate to **Policy Management > Gateway Firewall**.
- c) Click the **All shared rules** tab.
- d) Click **Add rule** (Add a policy first if needed) and, in the **Services** column, click the **Edit** button.
- e) On the **Set services** dialog box, on the **Services** tab, select the **ICMP destination unreachable** service, and click **Apply**.
- f) Click the **Settings** icon for the newly added rule and, on the **Settings** dialog box, activate the **Logging** toggle.
- g) In the **Applied to** column, click the **Edit** icon.
- h) In the **Applied to** dialog box, select the target NSX tier-0 gateway and click **Apply**.
- i) On the **Gateway Firewall** page, click **Publish**.
- j) Repeat the procedure for the **ICMP mask replies** and **ICMP redirects** services.

NOTE

A rule can also be created under Gateway Specific Rules to meet this requirement.

10. VMW-NSX-01532 NSX Tier-1 Gateway Firewall must be configured to inspect traffic at the application layer.
- a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Policy Management > Gateway Firewall** and select **Gateway Specific Rules**.
 - c) From the Gateway drop down choose **Tier-1 Gateway**

- d) For each rule that should have a Context Profile enabled, click the pencil icon in the **Profiles** column. and select **Context profile** under **Select profile** dialog box.
- e) Select an existing Context Profile or create a custom one then click **Apply**.
- f) After all the changes are made, click **Publish**.

NOTE

Not all App IDs will be suitable for use in all cases and should be evaluated in each environment before use.

A list of App IDs for application layer rules is available here: <https://docs.vmware.com/en/NSX-Application-IDs/index.html>

11. VMW-NSX-01469 Unicast Reverse Path Forwarding (uRPF) must be enabled on the NSX Tier-0 Gateway
 - a) On the main navigation bar, click **Networking**.
 - b) In the left pane, navigate to **Connectivity › Tier-0 gateways**.
 - c) Expand the NSX tier-0 gateway.
 - d) Expand the **Interfaces and GRE Tunnels** section and click the number of **External and Service Interfaces**.
 - e) In the **Set Interfaces** dialog box, click the vertical ellipsis and click **Edit** for the first interface.
 - f) From the drop-down set the **URPF Mode** to **Strict** and then click **Save**.
 - g) Repeat the step to configure all interfaces.
12. VMW-NSX-01459 VMW-NSX-01470 The NSX Tier-0 Gateway router must be configured to use encryption for BGP routing protocol authentication and use a unique password for each autonomous system (AS) that it peers with.
 - a) On the main navigation bar, click **Networking**.
 - b) In the left pane, navigate to **Connectivity › Tier-0 gateways**.
 - c) Expand the NSX tier-0 gateway.
 - d) Expand the **BGP** section and click the number next to **BGP neighbors**.
 - e) In the **Set BGP neighbors** dialog box, click the vertical ellipsis and click **Edit** for the first neighbor.
 - f) Under **Timers and Password**, enter a unique password of up to 20 characters that is different from other autonomous systems and then click **Save**.
 - g) Repeat the step to configure all neighbors.
13. VMW-NSX-01536 The NSX Tier-0 Gateway router must be configured to use encryption for OSPF routing protocol authentication.
 - a) On the main navigation bar, click **Networking**.
 - b) In the left pane, navigate to **Connectivity › Tier-0 gateways**.
 - c) Expand the NSX tier-0 gateway.
 - d) Expand the **OSPF** section and click number next to **Area Definition**.
 - e) In the **Set Area Definition** dialog box, click the vertical ellipsis and click **Edit** for the first Area definition.
 - f) Change the **Authentication** drop-down to MD5 and enter a Key ID and password and then click **Save**.
 - g) Repeat the step to configure all Area definitions.

NOTE

The MD5 password can have a maximum of 16 characters.

Configure Security Settings for NSX Edge Nodes by Using CLI Commands

You configure the NSX Gateway Firewall to send logs to a central log server.

You perform these procedures on the NSX tier-0 and tier-1 gateway only if your environment uses NSX Edges.

1. In a Web browser, log in to vCenter Server by using the vSphere Client.

Setting	Value
URL	<code>https://management-domain-vcenter-server-fqdn/ui</code>
User name	<code>administrator@vsphere.local</code>

2. In the **VMs and templates** inventory, navigate to the NSX Edge node, right-click the appliance, and select **Open remote console**.
3. VMW-NSX-01430, VMW-NSX-01511 Configure the NSX Gateway Firewall on the tier-0 and tier-1 gateways to send logs to a central log server.

You can configure the logging server with the LI-TLS or TLS protocols. You must store the server and client certificates in the `/var/vmware/nsx/file-store/` on each NSX Edge appliance.

- a) If you want to configure a TCP syslog server, run the command.

```
set logging-server <server-ip or server-name> proto tcp level info
```

- b) If you want to configure a TLS syslog server, run the command.

```
set logging-server <server-ip /_server-FQDN> proto tls level info serverca
ca.pem clientca ca.pem certificate cert.pem key key.pem
```

- c) If you want to configure a LI-TLS syslog server, run the command.

```
set logging-server <server-ip /_server-FQDN> proto li-tls level info serverca
root-ca.crt
```

NOTE

Configure the syslog or SNMP server to send an alert if the events server is unable to receive events from the NSX Edge node and if DoS incidents are detected.

Configure Security Settings for Distributed Firewall by Using the User Interface

You perform the procedure in NSX to configure traffic logging for Distributed Firewall rules, deny traffic by default and profiles such as spoof guard, flood protection, context, and IP Discovery profile.

1. In a Web browser, log in to the NSX Manager cluster as an administrator by using the user interface.
2. VMW-NSX-01409 Configure the NSX Distributed Firewall to generate traffic log entries.
 - a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Policy Management > Distributed Firewall**.
 - c) Click the **Category Specific Rules** tab.
 - d) For each rule with logging disabled, click the gear icon, activate the **Logging** toggle, and click **Apply**.
 - e) On the **Distributed Firewall** page, click **Publish**.
3. VMW-NSX-01412 Configure the NSX Distributed Firewall to deny network traffic by default and allow network traffic by exception.
 - a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Policy Management > Distributed Firewall**.
 - c) Click the **Category specific rules** tab and select **Application**.

- d) Expand the **Default Layer3 Section** and for the **Default Layer3 Rule**, select **Reject** or **Drop** from the **Actions** drop-down menu.
- e) On the **Distributed Firewall** page, click **Publish**.

CAUTION

Before denying, ensure the necessary rules to whitelist approved traffic are created and published or this change may result in a loss of communication for workloads.

4. VMW-NSX-01452 Configure flood protection profiles on the NSX Distributed Firewall to protect against Denial of Service (DDoS) attacks.
 - a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Settings > General Settings**.
 - c) Click the **Firewall > Flood Protection**.
 - d) From the **Add profile** drop-down menu, select **Add Firewall Profile**.
 - e) Enter a name and specify appropriate values for the following: **TCP half open connection limit**, **UDP active flow limit**, **ICMP active flow limit**, and **Other active connection limit**.
 - f) Enable **SYN Cache** and **RST Spoofing**.
 - g) Configure the **Applied to** field to contain appropriate security groups, and then click **Save**.
5. VMW-NSX-01534 The NSX Distributed Firewall must be configured to inspect traffic at the application layer.
 - a) On the main navigation bar, click **Security**.
 - b) In the left pane, navigate to **Policy Management > Distributed Firewall** and select **Category Specific Rules**.
 - c) For each rule that should have a Context Profile enabled, click the pencil icon in the **Profiles** column. and select **Context profile** under **Select profile** dialog box.
 - d) Select an existing Context Profile or create a custom one then click **Apply**.
 - e) After all the changes are made, click **Publish**.

NOTE

This control does not apply to ethernet rules.

Not all App IDs will be suitable for use in all cases and should be evaluated in each environment before use.

A list of App IDs for application layer rules is available here: <https://docs.vmware.com/en/NSX-Application-IDs/index.html>

Security Configurations Not Applicable or Not Compatible with VMware Cloud Foundation

Typical configuration guidelines apply to standalone implementations of VMware products. When these products are part of VMware Cloud Foundation, some configurations might not be applicable or might not be compatible with VMware Cloud Foundation. Do not implement these configurations. You can find mitigation steps for the configurations in the *VMware Cloud Foundation Audit Guide Appendix*.

Product	Configuration	Context for Excluding Configuration
vCenter Server	vCenter Server must be isolated from the public Internet but must still allow for patch notifications and delivery. VMW-VC-01231	Never apply patches to vCenter Server manually, using VMware vSphere Update Manager, or VMware vCenter Lifecycle Manager in a VMware Cloud Foundation environment, unless

Table continued on next page

Continued from previous page

Product	Configuration	Context for Excluding Configuration
		directed to do so by support. Patching the environment without using SDDC Manager might cause problems with automated upgrades or actions in the future.
ESXi	ESXi hosts using Host Profiles and/or Auto Deploy must use the vSphere Authentication Proxy to protect passwords when adding themselves to Active Directory. VMW-ESXI-00115	VMware Cloud Foundation does not use host profiles to join ESXi hosts to Active Directory.

Documentation Legal Notice

Information about the documentation legal notice.

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the "Documentation") is for your informational purposes only and is subject to change or withdrawal by Broadcom at any time. This Documentation is proprietary information of Broadcom and may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of Broadcom.

If you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all Broadcom copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to Broadcom that all copies and partial copies of the Documentation have been returned to Broadcom or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, BROADCOM PROVIDES THIS DOCUMENTATION "AS IS" WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL BROADCOM BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF BROADCOM IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice

The manufacturer of this Documentation is Broadcom Inc.

Provided with "Restricted Rights." Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b) (3), as applicable, or their successors.

Copyright © 2005–2025 Broadcom. All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies.

