

# VMware Cloud Foundation Operations and Administration Guide

23 JUN 2020

VMware Cloud Foundation 4.0

You can find the most up-to-date technical documentation on the VMware website at:

<https://docs.vmware.com/>

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2015-2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

About the VMware Cloud Foundation Operations and Administration Guide	7
<b>1 Administering Cloud Foundation Systems</b>	<b>8</b>
VMware Software Components Deployed in a Typical Cloud Foundation System	9
Web Interfaces Used When Administering Your Cloud Foundation System	9
<b>2 Getting Started with SDDC Manager</b>	<b>11</b>
Log in to the SDDC Manager Dashboard	11
Tour of the SDDC Manager User Interface	12
Log out of the SDDC Manager Dashboard	14
<b>3 Configuring Customer Experience Improvement Program</b>	<b>16</b>
<b>4 Certificate Management</b>	<b>18</b>
View Certificate Information	19
Configure a Microsoft Certificate Authority	19
Add OpenSSL CA support	21
Install Certificates with the Microsoft Certificate Authority or OpenSSL Certificate Authority	21
Install Certificates with External or Third-Party Certificate Authorities	24
Clean Out Old or Unused Certificates	28
<b>5 License Management</b>	<b>29</b>
Add License Keys for the Software in Your Cloud Foundation System	29
Edit License Description	30
Delete License Key	30
<b>6 Installing ESXi Software on Cloud Foundation Servers</b>	<b>31</b>
Download ESXi Software and VIBs	32
Provide Network Information for Imaging	32
Upload ISOs and VIBs to the VMware Imaging Appliance service	33
Image Servers with ESXi and VIBs	35
Post-Imaging Tasks	36
<b>7 Host Management</b>	<b>38</b>
Network Pool Management	38
Sizing a Network Pool	39
Create a Network Pool	40
View Network Pool Details	41

Edit a Network Pool	41
Delete a Network Pool	42
Enable Additional pNICs on Hosts	42
Commission Hosts	42
Decommission Hosts	46
View Host Inventory	47

## 8 Cluster Image Management 49

Create a Cluster Image	51
Making a Cluster Image Available in Cloud Foundation	52
Importing a Cluster Image	52
Extract a Cluster Image	55
View Available Cluster Images	55
Firmware Updates	56

## 9 Working with the Management Domain and VI Workload Domains 57

Adding Virtual Machines to the Management Domain	59
Storage for the Management Domain and VI Workload Domains	60
About VI Workload Domains	62
Prerequisites for a Workload Domain	63
Start the VI Configuration Wizard	65
Specify Names and Choose an Update Manager	65
Specify vSphere Cluster Details	66
Specify Compute Details	66
Specify Networking Details	66
Select the vSAN Parameters	67
Specify the VMFS on FC Datastore	68
Select Hosts	68
Specify NFS Storage Details	69
Select Licenses	69
View Object Names	70
Review Details and Start the Creation Workflow	70
Deploying NSX-T Edge Clusters	71
Create an NSX-T Edge Cluster	72
View Workload Domain Details	77
Delete a VI Workload Domain	77
View vSphere Cluster Details	78
Shrink a Workload Domain	79
Remove a Host from a vSphere Cluster in a Workload Domain	79
Delete a vSphere Cluster from a Workload Domain	80
Expand a Workload Domain	81

	Add a Host to a vSphere Cluster in a Workload Domain	81
	Add a vSphere Cluster to a Workload Domain	83
<b>10</b>	<b>Working with Workload Management</b>	<b>86</b>
	Sizing Compute and Storage Resources for a vSphere with Kubernetes Workload Domain	86
	Deploy Workload Management	87
	Configure NSX Route Maps on Edge T0 Router	88
	View Workload Management Cluster Details	90
<b>11</b>	<b>Deploy vRealize Suite Lifecycle Manager in Cloud Foundation</b>	<b>91</b>
<b>12</b>	<b>Download an Install Bundle</b>	<b>93</b>
<b>13</b>	<b>Multi-Instance Management</b>	<b>94</b>
	About the Multi-Instance Management Dashboard	96
	Create a Federation	99
	Invite a Cloud Foundation Instance to Join a Federation	100
	Join a Federation	101
	Join a Federation by Clicking an Invitation	102
	Join a Federation through the Multi-Instance Management Dashboard	103
	Leave a Federation	103
	Dismantle a Federation	104
<b>14</b>	<b>Stretching Clusters</b>	<b>105</b>
	About Availability Zones and Regions	105
	Stretched Cluster Requirements	106
	Deploy and Configure vSAN Witness Host	108
	Stretch a Cluster	109
	Expand a Stretched Cluster	115
	Unstretch a Cluster	118
	Replace a Failed Host in a Stretched Cluster	119
<b>15</b>	<b>Composability Management</b>	<b>122</b>
	Configure Translation Layer	122
	Compose a Server	124
	View Composability Information	125
	Add Storage	125
	Remove Storage	125
	Decompose a Server	126
<b>16</b>	<b>Monitoring Capabilities in the Cloud Foundation System</b>	<b>127</b>

Viewing Tasks and Task Details	127
<b>17</b>	<b>Updating Cloud Foundation DNS and NTP Servers 129</b>
Update DNS Server Configuration	129
Update NTP Server Configuration	132
<b>18</b>	<b>Supportability and Serviceability (SoS) Utility 136</b>
SoS Utility Options	136
Collect Logs for Your Cloud Foundation System	141
Component Log Files Collected by the SoS Utility	143
<b>19</b>	<b>Replacing Host Components 145</b>
Avoiding Unintentional Downtime	145
Replacing Components of a Host Running in Degraded Mode	146
Replace Components of an Assigned Host Running in Degraded Mode	146
Replace Components of an Unassigned Host Running in Degraded Mode	147
Replace a Dead Host	147
Replace Boot Disk on a Host	148
<b>20</b>	<b>User and Group Management 149</b>
Add a User or Group to Cloud Foundation	149
Remove a User or Group	150
Create a Service Account and Generate an Access Token	150
<b>21</b>	<b>Password Management 154</b>
Rotate Passwords	154
Manually Update Passwords	156
Look Up Account Credentials	157
Updating SDDC Manager Passwords	158
Update SDDC Manager Root and Super User Passwords	158
Update SDDC Manager REST API Account Password	159
Update Expired SDDC Manager root Password	159
<b>22</b>	<b>Backing Up and Restoring SDDC Manager and NSX Manager 161</b>
Image-Based Backup and Restore	162
File-Based Backup and Restore	162
Configure an External SFTP Server for File-Based Backups	163
Configure a Backup Schedule for SDDC Manager VM	164
Restore SDDC Manager	166
<b>23</b>	<b>Cloud Foundation Glossary 167</b>

# About the VMware Cloud Foundation Operations and Administration Guide

The *VMware Cloud Foundation Operations and Administration Guide* provides information about managing a VMware Cloud Foundation™ system, including managing the system's virtual infrastructure, managing users, configuring and deploying service offerings, and upgrading and monitoring the system.

## Intended Audience

The *VMware Cloud Foundation Operations and Administration Guide* is intended for cloud architects, infrastructure administrators, and cloud administrators who are familiar with and want to use VMware software to quickly deploy and manage an 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, virtual infrastructure (VI), and virtual desktop infrastructure (VDI)
- VMware virtualization technologies, such as VMware ESXi™, the hypervisor
- Software-defined networking using VMware NSX-T™ Data Center
- Software-defined storage using VMware vSAN™
- IP networks

Additionally, you should be familiar with these VMware software products, software components, and their features:

## Related Publications

The *Introducing VMware Cloud Foundation* document provides a high-level overview of the Cloud Foundation 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 Deployment Guide* contains detailed information about a Cloud Foundation system, its components, and the network topology of a deployed system.

The *VMware Cloud Foundation Lifecycle Management* document describes how to manage the lifecycle of a Cloud Foundation environment.

# Administering Cloud Foundation Systems

# 1

As an SDDC administrator, you use the information in the *VMware Cloud Foundation Operations and Administration* document to understand how to administer and operate your installed Cloud Foundation system.

An administrator of a Cloud Foundation system performs tasks such as:

- Manage certificates and passwords.
- Add capacity to your system.
- Configure and provision the systems and the workload domains that are used to provide service offerings.
- 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 Cloud Foundation software components.

---

**Note** Perform all Cloud Foundation operations in the SDDC Manager UI. Do not use the vSphere Client/Web Client or VMware Host Client to modify or delete resources which Cloud Foundation has deployed and configured, unless specifically instructed to do so in the Cloud Foundation documentation.

---

See the *Introducing VMware Cloud Foundation* document for a high-level overview of the Cloud Foundation product and the *VMware Cloud Foundation Deployment Guide* for information on deploying the product.

This chapter includes the following topics:

- [VMware Software Components Deployed in a Typical Cloud Foundation System](#)
- [Web Interfaces Used When Administering Your Cloud Foundation System](#)



# VMware Software Components Deployed in a Typical Cloud Foundation System

In a typical Cloud Foundation system, you encounter specific VMware software that SDDC Manager deploys in the system.

---

**Note** For information about which specific editions of each VMware product are licensed for use with the Cloud Foundation license, use the information resources at the Cloud Foundation product information page at <http://www.vmware.com/products/cloud-foundation.html>.

---

For the exact version numbers of the VMware products that you will see in your Cloud Foundation system after the initial bring-up process, see the *Release Notes* document for your Cloud Foundation version. If the system has been updated after the initial bring-up process using the Life Cycle Management features, see "View Upgrade History" in the *VMware Cloud Foundation Upgrade Guide* for details on how to view the versions of the VMware software components that are within your system.

---

**Caution** Do not manually change any of the settings that SDDC Manager sets automatically. If you change the generated settings, like names of VMs, port groups, virtual switches, or resource pools, unpredictable results might occur. Do not change settings for the resources that are automatically created and deployed during workflows, the workload domain processes, assigned IP addresses or names, and so on.

---

You can find the documentation for the following VMware software products and components at [docs.vmware.com](https://docs.vmware.com):

- vSphere (vCenter Server and ESXi)
- vSAN
- NSX-T Data Center
- vRealize Suite

## Web Interfaces Used When Administering Your Cloud Foundation System

You use SDDC Manager loaded in a browser for the single-point-of-control management of your Cloud Foundation system. 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 Dashboard, you can use the following user interfaces for administration tasks involving their associated VMware software components that are part of a VMware SDDC. All these interfaces run in a browser, and you can launch them from within the SDDC Manager Dashboard.

Launch links are typically identified in the user interface by the launch icon: .

VMware SDDC Web Interfaces	Description	Launch Link Location in SDDC Manager Dashboard
vSphere Client	This interface provides direct management of resources managed by the vCenter Server instances, for identity management, and for management of the NSX resources that provide the software-defined networking capabilities of the SDDC. You can also manage object level storage policies for distributed software-defined storage provided by vSAN.	<ol style="list-style-type: none"> <li>1 On the SDDC Manager Dashboard, click <b>Inventory &gt; Workload Domains</b>.</li> <li>2 Click <b>View Details</b> for a workload domain.</li> <li>3 In the Domain column, click the domain name.</li> <li>4 Click the <b>Services</b> tab.</li> <li>5 Click the vCenter Server launch link.</li> </ol>
NSX-T Manager Web interface		<ol style="list-style-type: none"> <li>1 On the SDDC Manager Dashboard, click <b>Inventory &gt; Workload Domains</b>.</li> <li>2 Click <b>View Details</b> for a workload domain.</li> <li>3 In the Domain column, click the domain name.</li> <li>4 Click the <b>Services</b> tab.</li> <li>5 Click the NSX-T Cluster launch link.</li> </ol>

# Getting Started with SDDC Manager

## 2

You use SDDC Manager to perform administration tasks on your Cloud Foundation system. This user interface 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 Dashboard by loading it in a web browser. For the list of supported browsers and versions, see the *Release Notes*.

---

**Note** When performing out-of-band (OOB) troubleshooting of hardware, some vendors may use Java-based consoles. Refer to the vendor documentation for supported browsers.

---

This chapter includes the following topics:

- [Log in to the SDDC Manager Dashboard](#)
- [Tour of the SDDC Manager User Interface](#)
- [Log out of the SDDC Manager Dashboard](#)

## Log in to the SDDC Manager Dashboard

You access SDDC Manager through the SDDC Manager Dashboard in a supported browser.

### Prerequisites


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.

### Procedure

- 1 In a browser, type one of the following:
  - `https://FQDN` where *FQDN* is the host name of the SDDC Manager.
  - `https://IP_address` where *IP\_address* is the IP address of the SDDC Manager.
- 2 Log in with the single-sign on user credentials.

### Results

You are logged in to SDDC Manager and the Dashboard page appears in the browser.



## Tour of the SDDC Manager User Interface

SDDC Manager provides the user interface for your single point of control for managing and monitoring your Cloud Foundation system and for provisioning virtual environments.

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 features and functions in the form of widgets. There are widgets for Solution; 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"> <li>■ To rearrange widgets, click the heading of the widget and drag it to the desired position.</li> <li>■ To hide a widget, hover the mouse anywhere over the widget to reveal the <b>X</b> in the upper-right corner, and click the <b>X</b>.</li> <li>■ To add a widget, click the three dots in the upper right corner of the page and select <b>Add New Widgets</b>. This displays all hidden widgets. Select a widget and click <b>Add</b>.</li> </ul>
Solutions	<p>Solutions includes the following section:</p> <ul style="list-style-type: none"> <li>■ Kubernetes - Workload Management enables you to start a Workload Management deployment and view Workload Management cluster details.</li> </ul>

Category	Functional Areas
Inventory	<p>Inventory includes the following sections:</p> <ul style="list-style-type: none"> <li>■ <b>Workload Domains</b> takes you to the Virtual Infrastructure page, which displays and provides access to all workload domains.</li> </ul> <p>This page includes detailed status and information about all existing workload domains, including IP addresses, health status, owner, number of hosts, and update status. It also displays CPU, memory, and storage utilization for each workload domain, and collectively across all domains.</p> <ul style="list-style-type: none"> <li>■ <b>Hosts</b> takes you to the Hosts page, which displays and provides access to current hosts and controls for managing hosts.</li> </ul> <p>This page includes detailed status and information about all existing hosts, including IP addresses, network pool, health status, domain and cluster assignment, and storage type. It also displays CPU, memory, and storage utilization for each host, and collectively across all hosts.</p>
Repository	<p>Repository includes the following sections:</p> <ul style="list-style-type: none"> <li>■ <b>Bundle Management</b> displays the available install, update, and upgrade bundles for your environment, as well as your bundle download history.</li> </ul> <hr/> <p><b>Note</b> To access bundles, you must be logged in to your My VMware account through the <b>Administration &gt; Repository Settings</b> page.</p> <hr/> <ul style="list-style-type: none"> <li>■ <b>Image Management</b> enables 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 lifecycle of ESXi hosts.</li> </ul>

Category	Functional Areas
Administration	<p>Administration includes the following sections:</p> <ul style="list-style-type: none"> <li>■ <b>Network Settings</b> enables you to configure, view, and manage network pool settings. You can create new network pools, and view and modify existing network pools.</li> <li>■ <b>Licensing</b> enables you to manage VMware product licenses. You can also add licenses for the component products in your Cloud Foundation deployment.</li> <li>■ <b>Users</b> enables you to manage Cloud Foundation users and groups, including creating users and groups, setting privileges, and assigning roles.</li> <li>■ <b>Repository Settings</b> enables you to log in to your My VMware account.</li> <li>■ <b>Composable Infrastructure</b> enables you to configure composable servers to meet the needs of your workloads without physically moving any hardware components.</li> <li>■ <b>vRealize Suite</b> enables you to deploy vRealize Suite Lifecycle Manager.</li> <li>■ <b>Security</b> enables you to configure your certificate authorities.</li> <li>■ <b>Backup Configuration</b> enables you to register an external SFTP server with SDDC Manager for backing up NSX Managers.</li> <li>■ <b>VMware CEIP</b> to enroll in the VMware Customer Improvement Plan. This plan provides VMware with information that enables 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 VMware products and services on a regular basis in association with your organization's VMware license key(s).</li> </ul>
Developer Center	<p>The Developer Center includes the following sections:</p> <ul style="list-style-type: none"> <li>■ <b>Overview:</b> API reference documentation. Includes information and steps for all the Public APIs supported by Cloud Foundation.</li> <li>■ <b>API Explorer:</b> Lists the APIs and allows you to invoke them directly on your Cloud Foundation system.</li> <li>■ <b>Code Samples:</b> Sample code to manage a Cloud Foundation instance.</li> </ul>

## Log out of the SDDC Manager Dashboard

Log out of SDDC Manager when you have completed your tasks.

## Procedure

- 1 In the SDDC Manager Dashboard, open the logged-in account menu by clicking the down arrow next to the account name in the upper right corner.
- 2 Click the menu choice to log out.

# Configuring Customer Experience Improvement Program


## 3

This product participates in VMware Customer Experience Improvement Program (CEIP).

The CEIP provides VMware with information that enables 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 the 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 turn CEIP on or off across all the SDDC 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 enable CEIP. Click **Apply**.



- You can turn CEIP on or off from the Administration tab on the SDDC Manager dashboard.

**Note** When you join CEIP, the **Enable VMware Customer Experience Improvement Program** task is shown in the floating task bar of the SDDC dashboard. Similarly, when you leave CEIP, the **Disable VMware Customer Experience Improvement Program** task is displayed. This task bar is available on all the screens of VMware Cloud Foundation.

To turn CEIP on or off from the **Administration** tab, perform the following steps:



## Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > VMware CEIP**.
- 2 To turn CEIP on, select the **Join the VMware Customer Experience Improve Program** option.
- 3 To turn CEIP off, deselect the **Join the VMware Customer Experience Improve Program** option.

# Certificate Management

# 4

You can manage certificates for all external-facing Cloud Foundation component resources, including configuring a certificate authority, generating and downloading CSRs, and installing them. This section provides instructions for using either the built-in OpenSSL Certificate Authority, which is part of SDDC Manager, or a Microsoft Certificate Authority.

You can manage the certificates for the following components.

- vCenter Server
- NSX Manager
- SDDC Manager
- vRealize Suite Lifecycle Manager

You replace certificates for the following reasons:

- Certificate has expired or is close to expiring.
- Certificate has been revoked.
- You do not want to use the default VMCA certificate.
- Optionally, when you create a new workload domain.

However, it is recommended that you replace all certificates right after deploying Cloud Foundation. After you create new workload domains, you can replace certificates for the appropriate components as needed.

## Procedure

### 1 [View Certificate Information](#)

You can view details of a currently active certificate for a component resource directly in the SDDC Manager Dashboard.

### 2 [Configure a Microsoft Certificate Authority](#)

Before you can generate and install certificates, you must configure a certificate authority (CA).

### 3 [Install Certificates with the Microsoft Certificate Authority or OpenSSL Certificate Authority](#)

You can generate a CSR and signed certificates and install them for selected resource components directly in the SDDC Manager Dashboard.

#### 4 Install Certificates with External or Third-Party Certificate Authorities

If you intend to generate and install external or third-party certificates, you must download the certificate signing request (CSR) from the SDDC Manager Dashboard and have it manually signed by a third-party CA. You can then use the controls in the SDDC Manager Dashboard to install the certificate.

#### 5 Clean Out Old or Unused Certificates

Old or unused certificates are stored in a trust store in SDDC Manager. You can delete old certificates through the SDDC Manager VM.

## View Certificate Information

You can view details of a currently active certificate for a component resource directly in the SDDC Manager Dashboard.

### Procedure

- 1 In the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 In the list of domains, click the name of the workload domain to open the details page for that domain.

The workload domain details page displays CPU, memory, and storage allocated to the domain.

- 3 Select the **Security Tab**.

This tab lists the certificates for each resource component associated with the workload domain. It displays the following details:

- Resource type
- Issuer, the certificate authority name
- Resource host name
- Valid from and valid to dates
- Certificate status: *Active*, *Expiring* (will expire within 15 days), or *Expired*.
- Certificate operation status.

- 4 To view certificate details, expand the resource to view the certificate details In the Resource Type column.

The expanded field displays certificate details including signature algorithm, public key, public key algorithm, certificate string, and more.

## Configure a Microsoft Certificate Authority

Before you can generate and install certificates, you must configure a certificate authority (CA).

## Prerequisites

- Verify that the Microsoft Certificate Authority Server has the correct roles 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 service 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.

---

**Note** If the CA Web server and CA are on different machines, you must perform the steps mentioned in <https://blogs.technet.microsoft.com/askds/2009/04/22/how-to-configure-the-windows-server-2008-ca-web-enrollment-proxy/> in addition to the following steps.

---

## Procedure

- 1 Navigate to **Administration > Security > Certificate Management** to open the Configure Certificate Authority page.
- 2 Click **Edit** and complete the following configuration settings.

Option	Description
Certificate Authority	Select the CA from the drop-down menu. The default is <b>Microsoft</b> .
CA Server URL	Specify the URL for the CA address server. This address must begin with <b>https://</b> and end with <b>certsrv</b> , for example <b>https://www.mymicrosoftca.com/certsrv</b>
Username	Provide a valid user name to enable access to the address server.
Password	Provide a valid password to enable access to the address server.
Template Name	Enter the certsrv template name. You must create this template in Microsoft Certificate Authority.

- 3 Click **Save**.

A dialog box appears, asking you to review and confirm the CA server certificate details.

- 4 Click **Accept** to complete the configuration.

## Results

The Microsoft CA is now available for use in generating and installing a certificate.

## Add OpenSSL CA support

To generate OpenSSL Certificate Authority (CA) signed certificates for the VMware Cloud Foundation environment:

### Procedure

- 1 To configure the OpenSSL CA settings before generating the certificates, navigate to **Administration > Security > Certificate Management**.
- 2 In the Configure Certificate Authority page, select **OpenSSL** for **Certificate Authority**. Provide the required information.

Attribute	Description
Common Name	Specify the FQDN of OpenSSL CA.
Organizational Unit	Use this field to differentiate between divisions within your organization with which this certificate is associated.
Organization	Specify the name under which your company is legally registered. 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 legally registered.

Click **Save**.

- 3 To generate the OpenSSL CA signed certificates, navigate to **Inventory > Workload Domains > Select Domain**.
- 4 Under the **Security** tab, click **Generate Signed Certificates**.
- 5 The **Generate Signed Certificates** pop-up appears. Select **OpenSSL** as the Certificate Authority.
- 6 Click **Generate Certificates**.

## Install Certificates with the Microsoft Certificate Authority or OpenSSL Certificate Authority

You can generate a CSR and signed certificates and install them for selected resource components directly in the SDDC Manager Dashboard.

## Prerequisites

- Verify that the bring-up process is complete and successful.
- Verify that you have configured the Certificate Authority, as described in [Configure a Microsoft Certificate Authority](#).

## Procedure

- 1 In the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 In the list of domains, click the name of the workload domain to open the details page for that domain.

The workload domain details page displays CPU, memory, and storage allocated to the domain.

- 3 Select the **Security Tab**.

This tab lists the default certificates, among other details, for the Cloud Foundation resource components. It also provides controls for working with certificates.

---

**Note** You can view the current certificate and key information for a component by clicking the down-arrow icon next to the name.

---

- 4 Generate the CSR.
  - a Use the check boxes to select the resource components for which you want to generate the CSR.
  - b Click **Generate CSRS**.

The Generate CSRs dialog box opens.

- c Configure the following settings for the CSR.

Option	Description
<b>Algorithm</b>	Select the key type for the certificate. RSA (the default) is typically used. The key type defines the encryption algorithm for communication between the hosts.
<b>Key Size</b>	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
<b>Email</b>	Optionally, enter a contact email address.
<b>Organizational Unit</b>	Use this field to differentiate between divisions within your organization with which this certificate is associated.
<b>Organization</b>	Type name under which your company is legally registered. The listed organization must be the legal registrant of the domain name in the certificate request.
<b>Locality</b>	Type the city or locality where your company is legally registered.
<b>State or Province Name</b>	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
<b>Country</b>	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d Click **Generate CSR**.

The Generate CSRs dialog box closes. The Security tab displays a status of `CSR Generation is in progress`. When the CSR generation completes, the **Generate Signed Certificates** button becomes active.

## 5 Generate the signed certificates.

- a Leave all the resource components selected.
- b Click **Generate Signed Certificates**.

The Generate Signed Certificates dialog box appears, listing the selected components.

- c For the Select Certificate Authority, select the desired authority, and click **Generate Certificate**.
- d If you are using Microsoft Certificate Authority, ensure the following:
  - The values of the `ExtendedKeyUsage (EKU)` and `SAN` fields in the server certificate must be identical to the values in the CSR.
  - The key usage extension in the server certificate must include a keyword for either `digitalSignature` OR `keyEncipherment`.
  - The key usage extension in the root CA certificate must include both `crlSign` and `keyCertSign` keywords.

The Generate Signed Certificates dialog box closes. The Security tab displays a status of `Certificates Generation is in progress`. When the certificate generation completes, the **Install Certificates** button becomes active.

## 6 Click **Install Certificates**.

The Security tab displays a status of `Certificates Installation is in progress`.

---

**Note** As installation completes, the Certificates Installation Status column for each selected resource component in the list changes to `Successful` with a green check mark.

---

**Important** If you selected SDDC Manager as one of the resource components, you must manually restart SDDC Manager services to reflect the new certificate and to establish a successful connection between Cloud Foundation services and other resources in the management domain.

---

## 7 Restart all services using the provided `sddcmanager_restart_services.sh` script.

To restart the service:

- a Using SSH, log in to the SDDC Manager VM with the following credentials:

User name: **vcf**

Password: use the password specified in the deployment parameter workbook.

- b Enter **su** to switch to the root user.
- c Run the following command:

```
sh /opt/vmware/vcf/operationsmanager/scripts/cli/sddcmanager_restart_services.sh
```

# Install Certificates with External or Third-Party Certificate Authorities

If you intend to generate and install external or third-party certificates, you must download the certificate signing request (CSR) from the SDDC Manager Dashboard and have it manually signed by a third-party CA. You can then use the controls in the SDDC Manager Dashboard to install the certificate.

## Prerequisites

Verify that you have configured and packaged your certificate authority configuration files in the form of a `<domain_name>.tar.gz` file. The contents of this archive must adhere to the following structure:

- The name of the top-level directory must exactly match the name of the domain as it appears in the list on the **Inventory > Workload Domains** page. For example, `MGMT`.
- The PEM-encoded root CA certificate chain file (`rootca.crt`) must reside inside this top-level directory.



The `rootca.crt` file contains a root certificate authority and can have *N* number of intermediate certificates. The file structure of the `rootca.crt` file must look like the following 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.

- This directory must contain one sub-directory for each component resource.

The name of each sub-directory must exactly match the resource hostname of a corresponding component as it appears in the Resource Hostname column in the **Workload Domains > Security** tab.

For example, `nsxManager.vrack.vsphere.local`, `vcenter-1.vrack.vsphere.local`, and so on.

- 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 **Workload Domains > Security** tab. The content of the `.crt` file must end with a newline character. All certificates including `rootca.crt` must be in UNIX file format.

For example, the `nsxManager.vrack.vsphere.local` sub-directory would contain the `nsxManager.vrack.vsphere.local.crt` file.

- Additional requirements for NSX-T certificates are listed below.
  - Server certificate (`NSXT_FQDN.crt`) must contain the `Basic Constraints` field with value `CA:FALSE`.
  - Root CA certificate chain file (`rootca.crt`), intermediate certificates, and root certificate must contain the `Basic Constraints` field with value `CA:TRUE`.
  - If the NSX-T 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 > Security** tab.

---

## Procedure

- 1 In the SDDC Manager Dashboard, navigate to **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 In the list of domains, click the name of the workload domain to open the details page for that domain.

The workload domain details page displays CPU, memory, and storage allocated to the domain.

- 3 Select the **Security Tab**.

This tab lists the default certificates, among other details, for the Cloud Foundation resource components. It also provides controls for working with certificates.

---

**Note** You can view the current certificate and key information for a component by clicking the down-arrow icon next to the name.

---

- 4 Generate the CSR.

- a Use the check boxes to select the resource components for which you want to generate the CSR.

- b Click **Generate CSR**.

The Generate CSRs dialog box opens.

- c Configure the following settings for the CSR.

Option	Description
<b>Algorithm</b>	Select the key type for the certificate. RSA (the default) is typically used. The key type defines the encryption algorithm for communication between the hosts.
<b>Key Size</b>	Select the key size (2048 bit, 3072 bit, or 4096 bit) from the drop-down menu.
<b>Email</b>	Optionally, enter a contact email address.
<b>Organizational Unit</b>	Use this field to differentiate between divisions within your organization with which this certificate is associated.
<b>Organization</b>	Type name under which your company is legally registered. The listed organization must be the legal registrant of the domain name in the certificate request.
<b>Locality</b>	Type the city or locality where your company is legally registered.
<b>State or Province Name</b>	Type the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
<b>Country</b>	Type the country name where your company is legally registered. This value must use the ISO 3166 country code.

- d Click **Generate CSR**.

The Generate CSRs dialog box closes. The Security tab displays a status of `CSR Generation is in progress`. When CSR generation is complete, the **Download CSR** button becomes active.

- 5 Click **Download CSR** to download and save the CSR files to the directory structure described in the Prerequisites section above.

- 6 External to the SDDC Manager Dashboard, complete the following tasks:

- a Verify that the different `.csr` files have successfully generated and are allocated in the required file structure.
- b Get the certificate requests signed.  
This will create the corresponding `.crt` files.
- c Verify that the newly acquired `.crt` files are correctly named and allocated in the required file structure.
- d Package the file structure as `<domain name>.tar.gz`. The `<domain name>` folder must include the `rootca.crt` file.

- 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.

After you select the file, the **Upload** button becomes active.

- 9 Click **Upload**.

When upload is complete, the **Install Certificate** button becomes active.

- 10 Click **Install Certificate**.

The Security tab displays a status of `Certificates Installation is in progress`.

---

**Note** As installation completes, the Certificates Installation Status column for the affected components in the list changes to `Successful` with a green check mark.

---



---

**Important** If you selected SDDC Manager as one of the resource components, you must manually restart SDDC Manager services to reflect the new certificate and to establish a successful connection between Cloud Foundation services and other resources in the management domain.

---

- 11 Restart all services using the provided `sddcmanager_restart_services.sh` script.

To restart the service:

- a Using SSH, log in to the SDDC Manager VM with the following credentials:

User name: **vcf**

Password: use the password specified in the deployment parameter workbook.

- b Enter **su** to switch to the root user.
- c Run the following command:

```
sh /opt/vmware/vcf/operationsmanager/scripts/cli/sddcmanager_restart_services.sh
```

## Clean Out Old or Unused Certificates

Old or unused certificates are stored in a trust store in SDDC Manager. You can delete old certificates through the SDDC Manager VM.

### Procedure

- 1 Using SSH, log in to the SDDC Manager VM with the following credentials:

Username: **vcf**

Password: use the password specified in the deployment parameter workbook.

- 2 Enter **su** to switch to the root user.
- 3 Change to the `/opt/vmware/vcf/operationsmanager/scripts/cli` directory.

```
cd /opt/vmware/vcf/operationsmanager/scripts/cli
```

- 4 From the `/opt/vmware/vcf/operationsmanager/scripts/cli` directory, use the following script and command to discover the names of the certificates in the trust store.

```
sddcmanager-ssl-util.sh -list
```

- 5 Using the name of the certificate, delete the old or unused certificate.

```
sddcmanager-ssl-util.sh -delete <certificate alias name from list>
```

- 6 (Optional) Clean out root certificates in VMware Endpoint Certificate Store from the Platform Services Controller node.

See [Explore Certificate Stores from the vSphere Client](#) in the vSphere product documentation.

# License Management

# 5

In the deployment parameter workbook that you completed before bring-up, you entered license keys for the following components:

- VMware vSphere
- VMware vSAN
- VMware NSX-T Data Center
- VMware vCenter Server

After bring-up, these license keys appear in the Licensing screen of the SDDC Manager Dashboard.

You must have adequate license units available before you create a VI workload domain, add a host to a cluster, or add a cluster to a workload domain. Add license keys as appropriate before you begin any of these tasks.

This chapter includes the following topics:

- [Add License Keys for the Software in Your Cloud Foundation System](#)
- [Edit License Description](#)
- [Delete License Key](#)

## Add License Keys for the Software in Your Cloud Foundation System

You can add licenses to the Cloud Foundation license inventory.

### Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Click **+ License Key**.
- 3 Select the product key for which you are entering a license key.
- 4 Type the license key.

- 5 Type a description for the license.

A description can help in identifying the license.

- 6 Click **Add**.

#### What to do next

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 License Description

If you have multiple license keys for a product, the description can help in identifying the license. For example, you may want to use one license for high-performance workload domains and the other license for regular workload domains.

#### Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Hover your mouse in the license row that you want to edit.  
A set of three dots appears on the left of the product name.
- 3 Click the dots and then click **Edit Description**.
- 4 On the Edit License Key Description window, edit the description as appropriate.
- 5 Click **Save**.

## Delete License Key

Deleting a license key removes the license from the Cloud Foundation license inventory. If the license has been applied to any workload domain, host, or cluster, the license is not removed from them, but it cannot be applied to new workload domains, hosts, or clusters.

#### Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Administration > Licensing**.
- 2 Hover your mouse in the license row that you want to edit.  
A set of three dots appears on the left of the product name.
- 3 Click the dots and then click **Remove Key**.
- 4 On the Remove Key dialog box, click **Remove**.

#### Results

The license is removed from the Cloud Foundation license inventory

# Installing ESXi Software on Cloud Foundation Servers

# 6

You can use the VMware Imaging Appliance service (VIA) included with the VMware Cloud Builder appliance to image servers for use in the management domain and VI workload domains. Alternatively, you can install ESXi manually. For the supported ESXi version, see the BOM section of the *VMware Cloud Foundation Release Notes*.

You can use VIA to image servers prior to adding them to Cloud Foundation as part of the host commissioning process. For information about imaging servers prior to bring-up, see the *VMware Cloud Foundation Deployment Guide*.

You must have access to the VMware Cloud Builder appliance to use the VMware Imaging Appliance service. If you deleted the Cloud Builder appliance after bring-up, you can redeploy it as described in "Deploy Cloud Foundation Builder Appliance" in the *VMware Cloud Foundation Deployment Guide*.

## Server Prerequisites

The servers that you image must meet certain prerequisites:

- PXE Boot is configured as primary boot option
- Install device is configured as the second boot option
- Legacy boot mode configured in BIOS

---

**Note** Although the VMware Imaging Appliance service does not support UEFI boot mode, Cloud Foundation does support it for servers configured outside of VIA.

---

- Servers are in the same L2 domain as the VMware Cloud Builder appliance
- Servers are reachable over an untagged VLAN/Network (VLAN ID 0)
- The VMware Cloud Builder appliance is deployed on an untagged VLAN/Network
- Server hardware/firmware should be configured for virtualization and vSAN and match the Cloud Foundation BOM as described in the Release Notes
- Physical hardware health status should be "healthy" without any errors
- Any onboard NICs are deactivated on the servers and only the two 10 GbE NICs reserved for use with Cloud Foundation are enabled in BIOS

The default root credentials for servers imaged with VIA are user **root**, password **EvoSddc!2016**.

This chapter includes the following topics:

- [Download ESXi Software and VIBs](#)
- [Provide Network Information for Imaging](#)
- [Upload ISOs and VIBs to the VMware Imaging Appliance service](#)
- [Image Servers with ESXi and VIBs](#)
- [Post-Imaging Tasks](#)

## Download ESXi Software and VIBs

In order to image your servers, you need to download an ESXi ISO and any vSphere Installation Bundles (VIBs) required to get the servers to a supported version of ESXi. See the BOM section of the VMware Cloud Foundation Release Notes for information about ESXi support.

You can download the ISO and VIBs from My VMware (<https://my.vmware.com>) to any location on the Windows machine that is connected to the VMware Cloud Builder appliance. Make sure to record the MD5 or SHA-1 checksums. You will need them when you upload the ISO/VIB to the VMware Imaging Appliance service.

## Provide Network Information for Imaging

You must provide the VMware Imaging Appliance service with certain network information specific to your environment before you can image your servers. This information is contained in the `via.properties` file on the VMware Cloud Builder appliance.

### Procedure

- 1 SSH into the VMware Cloud Builder appliance using the credentials specified when you deployed the VM. See "Deploy Cloud Foundation Builder Appliance" in the *VMware Cloud Foundation Deployment Guide*.
- 2 Type **su** to switch to the root user.
- 3 Navigate to the `/opt/vmware/evorack-imaging/config/` directory.



#### 4 Update the `via.properties` file with your network information.

- a If the VMware Cloud Builder appliance is using the `eth0` interface (default), then you do not need to modify any of the properties in Section A. If the VMware Cloud Builder appliance has multiple network interfaces and is not using `eth0`, you must update the following properties.

Property	Description
<code>via.network.interface</code>	Interface of the VMware Cloud Builder appliance configured in management network.
<code>via.web.url</code>	The IP address used to access the VMware Imaging Appliance service UI. Update this with the IP address of VMware Cloud Builder appliance in the management network.
<code>via.network.ifaceaddr</code>	Update this with the IP address of VMware Cloud Builder appliance in the management network.
<code>via.dhcp.esxi.tftpServer</code>	IP address of the server where TFTP is running. Update this with the IP address of VMware Cloud Builder appliance in the management network.
<code>via.config.remote.pxe=false</code>	Do not modify.

- b Update Section B with the network information for your environment.

Property	Description
<code>via.dhcp.netmask</code>	Netmask of the management network.
<code>via.dhcp.subnet</code>	Subnet of the management network.
<code>via.dhcp.routers</code>	Gateway IP of the management network.
<code>via.esxi.firewall.allowed.network</code>	CIDR notation for subnet IP of the management network.

#### 5 Type `systemctl restart imaging.service` to restart the imaging service.

Wait for the imaging service to restart.

#### 6 Type `systemctl status imaging.service` to verify that the imaging service is running.

#### What to do next

Log in to the VMware Imaging Appliance service and upload software.

## Upload ISOs and VIBs to the VMware Imaging Appliance service

After you have downloaded the required software and updated `via.properties` with your network information, you can upload ISOs and VIBs to the VMware Imaging Appliance service.

## Procedure

- 1 In a web browser on the Windows machine that is connected to the VMware Cloud Builder appliance, navigate to `https://Cloud_Builder_VM_IP:8445/via`.

The VMware Imaging Appliance service page displays.

- 2 Enter the admin credentials you provided when you deployed the VMware Cloud Builder appliance and click Log in.
- 3 Click **Bundle** and then click the **ESXi ISOs** tab.
- 4 Click **Browse** to locate and select the ISO.

The screenshot shows the VMware Imaging Appliance web interface. The top navigation bar includes 'VMware Imaging Appliance', 'Bundle' (selected), 'Imaging', 'History', 'Logs', 'About', a notification bell with '(0)', and 'Logout'. Below the navigation bar, the 'ESXi ISOs' tab is selected, and the 'Modify VIBs' sub-tab is active. The main content area shows a table for 'Available ISOs' with columns 'Name', 'Source', and 'Activate'. Below the table, there is a 'Select ISO to Add' field with the value 'VMware-VMvisor-installer-7.0.0-15799291.x86\_64.iso' and a 'Browse' button. There is also an 'MD5 Checksum' field, a 'Checksum Type' dropdown set to 'MD5' with 'SHA-1' as an option, and an 'ESXi License Key' field with the value 'Optional'. At the bottom, there is an 'Upload ISO' button.

- 5 Select the checksum type and enter the checksum.
- 6 Click **Upload ISO**.
- 7 When the uploaded ISO appears, select **Activate** to use the ISO for imaging servers.
- 8 Click the **Modify VIBs** tab.

The steps for uploading VIBs are optional.

- 9 Click **Browse** to locate and select the VIB.

The screenshot shows the VMware Imaging Appliance web interface with the 'Modify VIBs' sub-tab selected. The 'Available VIBs' table has columns 'Name' and 'In use'. Below the table, there is a 'Select VIB' field with the value 'VMware\_bootbank\_cpu-microcode\_7.0.0-1.0-15799291.vib' and a 'Browse' button. At the bottom, there is an 'Upload VIB' button.

- 10 Click **Upload VIB**.
- 11 When the uploaded VIB appears, select **In use** to use the VIB for imaging servers.

## What to do next

Use the selected ISO and VIB(s) to image servers for use with Cloud Foundation.

## Image Servers with ESXi and VIBs

Once you have uploaded the required ESXi and VIB packages to the VMware Imaging Appliance service, you can begin imaging servers. You can image an individual server, or multiple servers at the same time.

You can use VIA to image servers for use in the management domain and VI workload domains. The management domain requires a minimum of four servers. See the *Planning and Preparation Workbook* for more information about requirements.

**Note** When you image servers, VIA uses the ESXi ISO that you activated and the VIB(s) that you marked as **In use**.

### Procedure

- 1 In a web browser on the Windows machine that is connected to the VMware Cloud Builder appliance, navigate to `https://Cloud_Builder_VM_IP:8445/via`.

The VMware Imaging Appliance service page displays.

- 2 Enter the admin credentials you provided when you deployed the VMware Cloud Builder appliance and click Log in.
- 3 Click Imaging.
- 4 Enter the required information.

Name	<input type="text" value="MGMT Domain"/>
Description	<input type="text" value="Servers for the management domain"/>

NOTE : Please ensure that the user guidelines and pre requisites are followed as per product documentation before proceeding.

ESXI SERVER

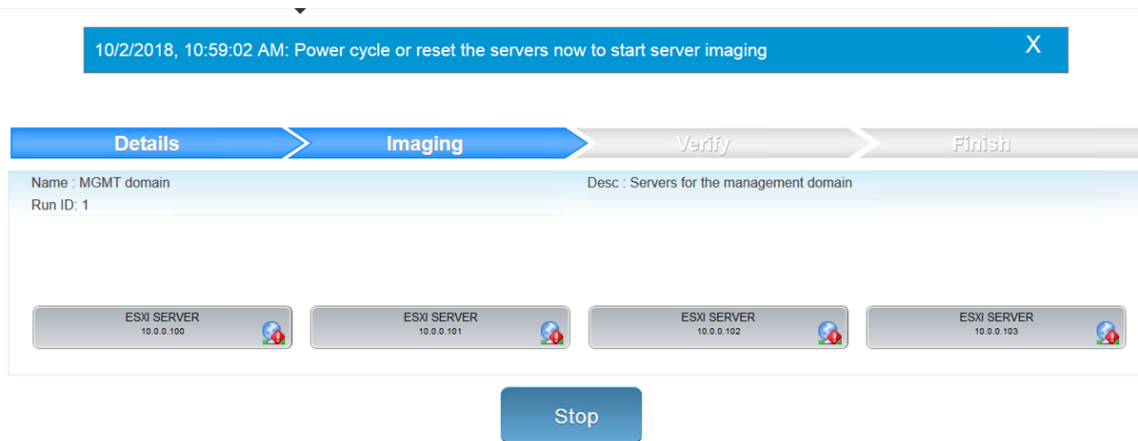
NTP Server:  Number:

①	IP: <input type="text" value="10.0.0.100"/>	MAC: <input type="text" value="00:50:56:ad:7a:1b"/>	Host FQDN: <input type="text" value="esxi-1.vrack.vsphere.local"/>
②	IP: <input type="text" value="10.0.0.101"/>	MAC: <input type="text" value="00:50:56:ad:75:26"/>	Host FQDN: <input type="text" value="esxi-2.vrack.vsphere.local"/>
③	IP: <input type="text" value="10.0.0.102"/>	MAC: <input type="text" value="00:50:56:ad:28:35"/>	Host FQDN: <input type="text" value="esxi-3.vrack.vsphere.local"/>
④	IP: <input type="text" value="10.0.0.103"/>	MAC: <input type="text" value="00:50:56:ad:24:f8"/>	Host FQDN: <input type="text" value="esxi-4.vrack.vsphere.local"/>

Option	Description
Name	Enter a name for the imaging job.
Description	Enter a description for the imaging job.
NTP Server	Enter the IP address for the NTP server.
Number	Enter the number of servers you want to image with the selected ISO and VIBs.

Option	Description
IP	Enter the IP address for the server.
MAC	Enter the MAC address for the server.
Host FQDN	Enter the FQDN for the server.

- 5 Click **Start Imaging**.
- 6 When prompted, power cycle the server(s) to continue imaging.



VIA displays information about the progress of imaging. Click a server to view details. Once imaging is complete, VIA performs verification of the servers.

- 7 When verification is finished, click **Complete**.

#### What to do next

Perform post-imaging tasks.

## Post-Imaging Tasks

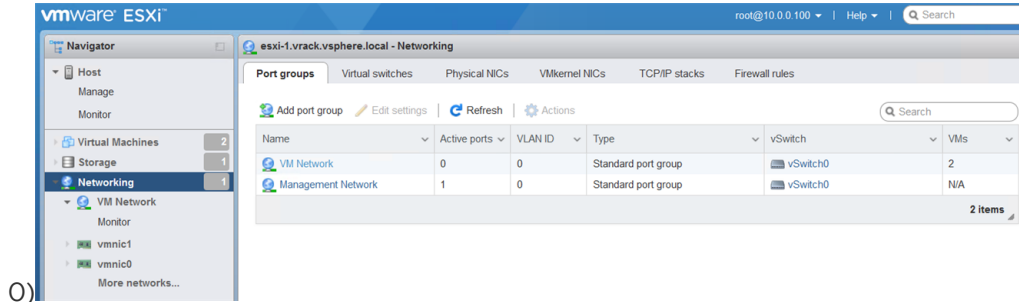
After you image your servers with ESXi and VIBs, you must perform some post-imaging tasks, depending on whether you use an untagged or a tagged management VLAN.

For imaging servers, the VMware Imaging Appliance service requires an untagged VLAN. You can continue to use an untagged VLAN for management, or you can use a tagged VLAN.

## Untagged Management VLAN

In this scenario, you use the same network for provisioning and management.

- Ensure that the Management Network and VM Network port groups on each host use the untagged VLAN (VLAN ID



- Verify that your DNS and NTP server are routable to the management network and ESXi hosts can reach them. To configure a default gateway or static routes on your ESXi hosts, see <https://kb.vmware.com/kb/2001426>.

## Tagged Management VLAN

In this scenario, you use an untagged VLAN for provisioning and a tagged VLAN for management.

- Modify the Management Network and VM Network port groups on each host to use the tagged VLAN
- Migrate the hosts from the provisioning network to the management network on the TOR switches
- Verify that your DNS and NTP server are routable to the management network and ESXi hosts can reach them. To configure a default gateway or static routes on your ESXi hosts, see <https://kb.vmware.com/kb/2001426>.

# Host Management

# 7

To add hosts to the Cloud Foundation 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 enable the additional pNICs before adding them to Cloud Foundation.

For information on network pools, see [Network Pool Management](#). For information on enabling additional pNICs, see [Enable Additional pNICs on Hosts](#).

You then commission hosts to Cloud Foundation. During the commissioning process, you associate hosts with a network pool. Commissioned hosts are added to the Cloud Foundation inventory. You can add these hosts to the management domain or to a VI workload domain. 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.

This chapter includes the following topics:

- [Network Pool Management](#)
- [Enable Additional pNICs on Hosts](#)
- [Commission Hosts](#)
- [Decommission Hosts](#)
- [View Host Inventory](#)

## Network Pool Management

When you create a VI workload domain or add a host or vSphere cluster to a workload domain, 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 a set of subnets within an L2 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 Cloud Foundation inventory.

**Table 7-1. Information Required for a Network Pool**

Storage Being Used	Required Networks in Network Pool
vSAN	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
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 vMotion and vSAN, NFS, or iSCSI 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 L2 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 workload domain with hosts in a different L2 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. For NSX-T workload domains, you can use the Cloud Foundation API to select hosts from different network pools, if those network pools have the same VLAN ID and MTU settings.

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 domains (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. Original servers used in the first cluster can get outdated at some point. Newer server models can then be added in a new cluster to the workload domain and workloads can be migrated at a leisurely pace.

## Sizing 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 workload domain must contain a minimum of one vSphere cluster, with a minimum number of three hosts. The exception to this rule is the management workload domain, which requires a minimum of four hosts. This 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.

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 L3 (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 L2 domain
- Number of hosts required in a vSphere cluster
- Whether the network switches will be shared with non-Cloud Foundation hosts
- Number of workload domains you plan on creating

## Create a Network Pool

A network pool must include vMotion network information. Depending on the type of storage you are using, you may also need to provide network information for vSAN and NFS.

The subnet in a network pool cannot overlap the subnet of another pool.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Network Settings**.
- 2 Click **Create Network Pool**.
- 3 Enter a name for the network pool.



#### 4 Select the network type(s).

You can include both vSAN and NFS network information in the same network pool or create separate network pools for vSAN and NFS. For VMFS on FC storage, select vMotion only or vMotion and NFS.

#### 5 Provide the following network information for the selected network type(s).

- a Enter a VLAN ID between 1 and 4094.
- b Enter an MTU between 1500 and 9000.
- c In the **Network** field, enter a subnet IP address.
- d Enter the subnet mask.
- e Enter the default gateway.
- f 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.

---

**Note** Ensure that you have entered the correct IP address range. IP ranges cannot be edited after the network pool is created.

---

#### 6 Click **Save**.

## View Network Pool Details

You can view vSAN and vMotion network details for a network pool as well as the total number of used and available IP addresses.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Network Settings**.
- 2 Click the arrow to the left of the pool name.

A high-level summary of the network pool's vSAN and vMotion network information is displayed.

- 3 Click the name of a network pool.

Network pool details are displayed.

## Edit a Network Pool

You can add an IP inclusion range to a network pool. No other parameters can be modified.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Network Settings**.

- 2 Hover your mouse in the network pool row that you want to edit.

A set of three dots appear on the left of the pool name. Click these dots and then click **Edit**.

- 3 Enter an IP inclusion range and click **Add**.
- 4 Click **Save**.

## Delete a Network Pool

You can delete a network pool if none of the hosts with an IP address from this pool belong to a workload domain. The default pool created during bring-up cannot be deleted.

### Prerequisites

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.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Network Settings**.
- 2 Hover your mouse in the network pool row that you want to delete.  
A set of three dots appear on the left of the pool name. Click these dots and then click **Delete**.

## Enable Additional pNICs on Hosts

You can isolate VMkernel traffic (management, vSAN, vMotion, or overlay) across multiple physical NICs (pNICs). Each vSphere distributed switch supports two pNICs, so you must define additional vSphere distributed switches and map traffic per vDS as appropriate.

Traffic isolation across physical NICs is not supported through the UI. You must enable the additional pNICs on hosts before commissioning them to Cloud Foundation. You can then update the API spec to map traffic flow to reflect your physical topology. For information on using APIs, see *VMware Cloud Foundation API Reference Guide*.

## Commission Hosts

Adding hosts to the Cloud Foundation inventory is called commissioning. You can add hosts individually or use a JSON template to add multiple hosts at once.

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, Cloud Foundation validates and commissions each host. Each host is added to the free pool and is available for workload domain creation.

The storage type you select for a host (vSAN, NFS, VMFS on FC), must be supported by its associated network pool. A network pool can support both vSAN and NFS. For VMFS on FC storage, the network pool must be vMotion only or vMotion and NFS.

- Hosts that use vSAN storage can only be used with vSAN-based workload domains.

- 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.

---

**Note** The management domain can only include hosts that use vSAN storage.

---

See [VMware Configuration Maximums](#) for information about the maximum number of hosts you can commission at one time.

### Prerequisites

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.
- Host has the drivers and firmware versions specified in the VMware Hardware Compatibility Guide.
- Hardware health status is healthy without any errors.
- A supported version of ESXi is installed on the host. See the *VMware Cloud Foundation Release Notes* for information about supported versions.
- 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 should be restricted to the management port group.

You can commission hosts with more than two NICs using the Cloud Foundation API.

- Management IP address is configured on the first NIC port.
- Host is configured with appropriate gateway. The gateway must be part of the management subnet.
- SSH and syslog are enabled.
- DNS is configured for forward and reverse lookup and FQDN.
- All disk partitions on HDD and SSD are deleted.

---

**Note** You must have a network pool available in order to commission a host.

---

### Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Hosts**.

- 2 Click **Commission Hosts**.
- 3 Confirm that hosts to be commissioned meet each criterion in the checklist and select the check boxes.
- 4 Click **Proceed**.

- 5 Select whether you want to add hosts one at a time or import a JSON file to add multiple hosts at once.

Option	Description
<b>Add new</b>	<p>Manually enter the following information for the host you want to add:</p> <ul style="list-style-type: none"> <li>■ FQDN</li> <li>■ Network pool (choose an existing network pool from the list)</li> <li>■ User name and password (root credentials)</li> <li>■ Storage type (vSAN, NFS, or VMFS on FC)</li> </ul> <p>Host Addition and Validation ⓘ</p> <hr/> <p>▼ Add Hosts</p> <p>You can either choose to add host one at a time or download <a href="#">JSON</a> template and perform bulk commission.</p> <p><input checked="" type="radio"/> Add new <input type="radio"/> Import</p> <p>Host FQDN <input type="text" value="esxi-5.vrack.vsphere.local"/></p> <p>Storage Type <input checked="" type="radio"/> vSAN <input type="radio"/> NFS <input type="radio"/> VMFS on FC</p> <p>Network Pool Name ⓘ <input type="text" value="bringup-networkpool"/></p> <p>User Name <input type="text" value="root"/></p> <p>Password <input type="password" value="....."/> ⓘ</p> <p><b>ADD</b></p>

Click **Add**.

You can now add more hosts or proceed to the next step.

#### Import


- Click the link to download the JSON template.
- Open the JSON template file and enter information about the hosts to add.
  - FQDN
  - User name and password (root credentials)
  - Storage type (vSAN, NFS, or FC)
  - Network pool name

```

1  {
2      "hostsSpec": [
3          {
4              "hostfqdn": "esxi-5.vrack.vsphere.local",
5              "username": "root",
6              "password": "Er5!x98b",
7              "storageType": "vSAN",
8              "networkPoolName": "bringup-networkpool"
9          },
10         {
11             "hostfqdn": "esxi-6.vrack.vsphere.local",
12             "username": "root",
13             "password": "B89x!5rE",
14             "storageType": "NFS",
15             "networkPoolName": "sfo-networkpool"
16         }
17     ]
18 }
  
```

Option	Description
	c Click <b>Browse</b> to locate and select the JSON file containing host information.
	d Click <b>Upload</b> .

The host or hosts appear in the **Hosts Added** section.

- 6 Verify that the server fingerprint is correct for each host and then click the confirm fingerprint icon .

- 7 Click **Validate All**.

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** to review the host information and then click **Commission** to begin commissioning.

The Hosts page appears, and the status of the commission task is displayed. Click **View Status in Task** to display the task bar.

### Results

The commissioned hosts are added to the host table. The host belongs to a free pool until you assign it to a workload domain.

## Decommission Hosts

Removing hosts from the Cloud Foundation inventory is called decommissioning. You can decommission a host for maintenance work or if you want to add it to another network pool. If you want to re-use a host in a different workload domain, you must decommission the host and re-image it up before adding it to the workload domain.

See [VMware Configuration Maximums](#) for information about the maximum number of hosts you can decommission at one time.

### Prerequisites

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](#).

### Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Hosts**.
- 2 Click **Unassigned Hosts**.
- 3 In the hosts table, select the host(s) you want to decommission.
- 4 Click **Decommission Selected Hosts**.

## 5 Click **Confirm**.

The Hosts page appears and the status of the decommission task is displayed. Click **View Status in Task** to display the task bar.

### What to do next

Re-image the decommissioned host before adding it to a workload domain.

## View Host Inventory

The Hosts page displays details about all the hosts in your Cloud Foundation system, 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 name
- IP address
- The network pool to which the host belongs
- Current status
- Host state, or the workload domain to which it is allocated
- Cluster or more specifically, the domain cluster to which it is assigned
- Host-specific CPU and memory usage
- Storage type

The Hosts page also provides controls for commissioning hosts.

### Procedure

- 1 From the SDDC Manager Dashboard, navigate to **Inventory > Hosts**.

The Hosts page appears.

- 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 name 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 including capacity and cache tiers.

---

**Note** Below the page title, the host details page also provides quick links to the network pool and the workload domain cluster to which the host belongs.

---

- 4 (Optional) To decommission the host from the host details page, click **Actions** near the page name and select **Decommission**.

For details, see [Decommission Hosts](#).

- 5 (Optional) To view host VM details, click **Actions** near the page name and select **Open in ESXi Client**.

The ESXi Client opens.



# Cluster Image Management

## 8

In Cloud Foundation 4.0, you can use cluster images as an alternative way of performing ESXi host lifecycle operations. A cluster image represents a desired software specification to be applied to all hosts in a vSphere cluster. Software and firmware updates happen simultaneously, in a single workflow.

A cluster image is a precise description of the software, components, vendor add-ons, and firmware to run on a host. With this new functionality, you set up a single image and apply it to all hosts in a cluster, thus ensuring cluster-wide host image homogeneity.

## Cluster Image and vSphere Lifecycle Manager

Cluster images are made available by the vSphere Lifecycle Manager (vLCM), a vCenter service. This service is now integrated with Cloud Foundation and enables centralized and simplified lifecycle management of ESXi hosts. When a VI workload domain or cluster is created with an image, you can update and upgrade the ESXi software on all hosts in a cluster. You can also install driver add-ons, components, and firmware on the hosts.

For more information on cluster images (also referred to as vSphere Lifecycle Manager Images), see [vSphere Lifecycle Manager Images](#).

## Cluster Image Components

A cluster image may consist of four elements: ESXi base image, a vendor add-on, a firmware and drivers add-on, and additional components. It is mandatory to add ESXi to a cluster image. Adding the other elements is optional.

- ESXi base image

The base image is an ESXi image that VMware provides with every release of ESXi. The base image is a collection of components that is complete and can boot up a server. Base images have a user-readable name and a unique version.

- Vendor Add-On

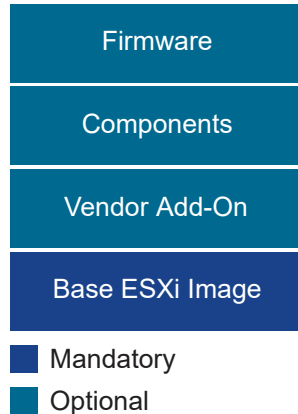
A vendor add-on is a collection of software components for the ESXi hosts that OEMs create and distribute. This vendor add-on can contain drivers, patches, and solutions.

- 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. The firmware and drivers add-on 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.



## Cluster Image in Cloud Foundation

Cluster images must be created in vSphere 7.0 and then imported to Cloud Foundation. Unlike vSphere where cluster images are managed per cluster, Cloud Foundation allows you to manage all cluster 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 Cloud Foundation. Since an image is created for a cluster of hosts, you must create a vLCM-enabled cluster 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.

After the image is created in vSphere, you export the cluster image specification and component files from vSphere to your local computer. You then import these files to Cloud Foundation so that you can apply the image to a vLCM-enabled VI workload domain.

After you import an image to 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-T 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.

## Cluster Image Workflow



## Cluster Images are Optional for Workload Domains

You have two update manager options for VI workload domains. You select an update manager while creating a workload domain. All clusters in the workload domain are managed by the selected update manager.

- vLCM (vSphere Lifecycle Manager) automates the lifecycle management of your Cloud Foundation environment including firmware by using cluster images.
- vSphere Update Manager (VUM) also automates the lifecycle management of your Cloud Foundation environment but firmware updates are manual. The management domain is VUM-based.

This chapter includes the following topics:

- [Create a Cluster Image](#)
- [Making a Cluster Image Available in Cloud Foundation](#)
- [View Available Cluster Images](#)
- [Firmware Updates](#)

## Create a Cluster Image

Cluster images are created in vSphere 7.0. You can create an image either on the management domain vCenter Server, or a vCenter Server external to Cloud Foundation.

You first create an empty cluster in vSphere and then set up an image on that cluster. During the creation of an image, you define the ESXi version and can optionally add vendor add-ons, components, and firmware.

---

**Note** The default cluster in the management domain is not vLCM-enabled, so you cannot use this for creating an image.

---

### Prerequisites

If you want to add firmware to the cluster image, you must install the Hardware Support Manager from your vendor. See [Firmware Updates](#).

## Procedure

- 1 Look up the ESXi version supported by your Cloud Foundation release by referring to the BOM section in the *VMware Cloud Foundation Release Notes*.
- 2 Download the ESXi base image for this version and upload it to vSphere Client. See *VMware ESXi Installation and Setup*.
- 3 Create an empty cluster in the vCenter Server where you want to create a cluster image. You do not need to add any hosts to this cluster.
  - a In the vSphere Client, select **Actions > New Cluster**.
  - b As a best practice, name this cluster *ClusterForImage*.  
You can keep this cluster for creating additional cluster images for upgrades.
  - c Select **Manage all hosts in this cluster with a single image**.
  - d In the **Image Setup** section, select the ESXi software you uploaded in step 2. Optionally, you can also select a vendor add-on and components.
  - e If you have the vendor Hardware Support Manager installed, you can select firmware. This step is optional.
  - f Click **OK**.  
The new cluster and cluster image are created.

## What to do next

Import the cluster image to Cloud Foundation. See [Importing a Cluster Image](#) .

# Making a Cluster Image Available in Cloud Foundation

A cluster image must be made available in Cloud Foundation before it can be applied to a VI workload domain or cluster. You can either import a cluster image from vSphere 7.0 or extract an image from a vLCM-enabled workload domain cluster in SDDC Manager.

## Importing a Cluster Image

You export a cluster image from vSphere to your local computer and then import it to SDDC Manager.

The following files need to be uploaded to SDDC Manager.

Format	Content
JSON	Image specification
JSON	Cluster settings

Format	Content
ISO	ESXi image (optional) You can use the ISO file for imaging additional hosts that you bring into Cloud Foundation.
ZIP	Image components. The image specification is not included in this format.

### Prerequisites

A cluster image must have been created in vSphere 7.0. For more information, see [Create a Cluster Image](#).

### Procedure

- 1 Export the cluster image specifications and components from vSphere to your local computer.
  - a In the vSphere Client, click **Host and Clusters** and select the cluster from where you want to export the image.
  - b On the Updates tab, click **Image**.
  - c Click the horizontal ellipsis icon and select **Export**.
  - d Click **JSON** and then click **Export**.

This downloads the image specification JSON file.


- e 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 Cloud Foundation. Importing this file is optional.

---

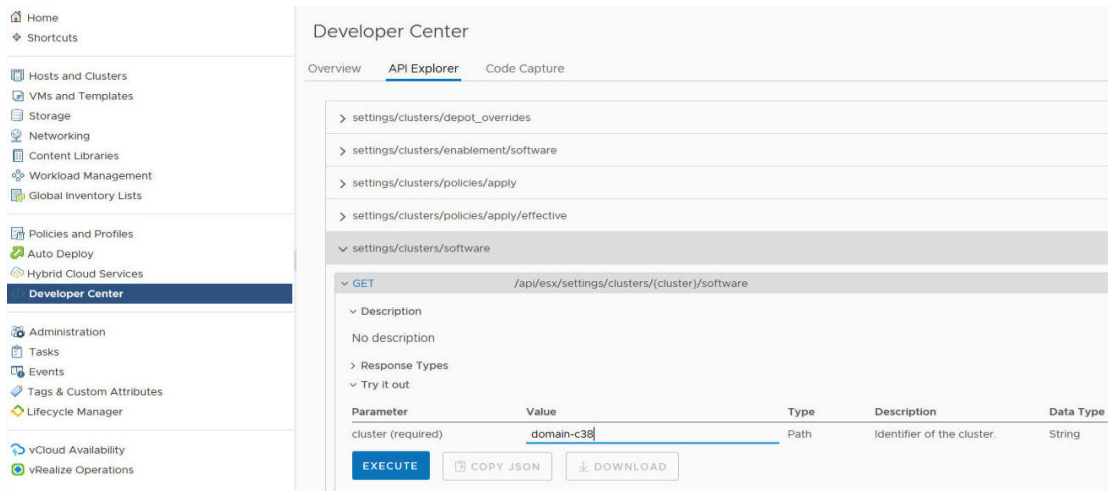
**Note** Do not rename these files.

---

- f Follow the steps below to download the cluster settings JSON file.
  - 1 From the Menu drop-down on the top bar, select **Developer Center**.
  - 2 Click **API Explorer** tab and ensure that the endpoint containing the cluster image is selected.
  - 3 Expand the cluster section.



- 4 Retrieve the cluster ID.
  - a Expand GET API for `/rest/vcenter/cluster` and click **Execute**.
  - b Scroll down to the Response area, expand the `vcenter.cluster.list_resp` link and then click the `vcenter.cluster.summary` link for the cluster that contains the image you want to import.
  - c Copy the value of the cluster parameter.
- 5 Scroll to the top of the Developer Center and in the **Select API** field, select `esx`.
- 6 Expand the `settings/clusters/software` API and then expand the GET for `/api/esx/settings/clusters/{cluster}/software` API.
- 7 In the value field for the cluster parameter, paste the value you had copied earlier.



- 8 Click **Execute** and then click **Download**.

The cluster settings JSON file (named `response-body.JSON`) is downloaded to your local computer.

- 2 In SDDC Manager, click **Repository > Image Management**.

- 3 In the Option 1 Import Cluster Image section, select the two JSON, zip, and ISO files from your local computer.

Ensure that you upload the image specification JSON and cluster setting JSON files in the correct fields. Mixing these files up will result in a validation error.

- 4 Enter a name for the cluster image and then click **Upload Image Components**.

#### Results

The 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 vLCM-enabled VI workload domain.

## Extract a Cluster Image

An image imported from vSphere can be customized when it is applied to a cluster in Cloud Foundation. For example, say your image included only the base ESXi image. When you use this image to create a VI workload domain, NSX-T 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.

#### Procedure

- 1 From the navigation bar, click **Repository > Image Management**.
- 2 In the Option 2 section, select a workload domain with an assigned image.
- 3 Select the cluster from where the image needs to be extracted.
- 4 Click **Import Cluster Image**.

#### Results

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 vLCM-enabled VI workload domain.

## View Available Cluster Images

You can view the cluster images available in Cloud Foundation.

#### Procedure

- 1 From the navigation bar, click **Repository > Image Management**.
- 2 Click the Available Images tab.

All cluster images available in your Cloud Foundation environment are displayed. You can search by image name, or filter by ESXi version or vendor.

## Firmware Updates

You can use cluster images to perform firmware updates on the ESXi hosts in a cluster. Using cluster images 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.

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.



# Working with the Management Domain and VI Workload Domains

# 9

The management domain and deployed VI workload domains are logical units that carve up the compute, network, and storage resources of the Cloud Foundation system. The logical units are groups of ESXi hosts managed by vCenter Server instances with specific characteristics for redundancy and VMware SDDC best practices.

The management domain is created by default during bring-up. The Cloud Foundation software stack is deployed on the management domain. Additional infrastructure virtual machines which provide common services, such as backup or security appliances, can be deployed in the management domain as well.

The management domain and workload domains include these VMware capabilities by default:

## **VMware vSphere<sup>®</sup> High Availability (HA)**

This feature supports distributed availability services for a group of ESXi hosts to provide rapid recovery from outages and cost-effective high availability for applications running in virtual machines. Out of the box, Cloud Foundation provides a highly available environment for workload domains. There may be additional settings (not set by default) that can increase availability even further. For more information about vSphere HA, see the *vSphere Availability* documentation at <https://docs.vmware.com/en/VMware-vSphere/>.

## **VMware vSphere<sup>®</sup> Distributed Resource Scheduler™ (DRS)**

This feature dynamically allocates and balances computing capacity across a group of hardware resources aggregated into logical resource pools or clusters. Clusters are the primary unit of operation in Cloud Foundation. DRS continuously monitors use across resource pools and allocates available resources among the virtual machines based on predefined rules that reflect business needs and changing priorities. When a virtual machine experiences an increased load, vSphere DRS automatically allocates additional resources by redistributing virtual machines among the physical servers in the resource pool. For more information about DRS, see the *vSphere Resource Management* documentation at <https://docs.vmware.com/en/VMware-vSphere/>.

## **VMware vSAN<sup>®</sup>**

This component aggregates local storage disks in a group of ESXi hosts to create a storage pool shared across all hosts in that group. For more information about vSAN, see the *VMware vSAN* documentation at <https://docs.vmware.com/en/VMware-vSAN/>.

## VMware NSX Manager Automated Backup

It is crucial to take backups of all NSX-T Data Center components to restore the system to its working state in the event of a failure. Until you register an external SFTP server, the NSX-T backups are taken on the SFTP server that is built into SDDC Manager. It is recommended that you register an external SFTP server soon after you upgrade or deploy VMware Cloud Foundation. See [Configure an External SFTP Server for File-Based Backups](#).

You can restore an NSX-T Data Center configuration back to the state that is captured in any of the backups.

Each Cloud Foundation instance is one SSO domain to which all vCenter Servers are joined. See [VMware Configuration Maximums](#) for information about the maximum number of supported workload domains.

---

**Note** if you use cross vCenter vMotion between two VI workload domains with dissimilar hardware, you must enable EVC on the corresponding clusters. See [Enable EVC on an Existing Cluster](#) in the vSphere product documentation. You can enable EVC on the management domain by selecting the appropriate value in the Deploy Parameters tab of the deployment parameter workbook. For more information, see the *VMware Cloud Foundation Deployment Guide*.

---

## Procedure

### 1 Adding Virtual Machines to the Management Domain

You can add virtual machines to the management domain as desired. Commonly, these virtual machines provide infrastructure services such as backup or security throughout the solution. To prevent resource conflicts between the core Cloud Foundation services, these additional virtual machines are added to the Compute-ResourcePool. This resource pool is automatically created during bring-up for this purpose.

### 2 Storage for the Management Domain and VI Workload Domains

To successfully create and manage a workload domain, Cloud Foundation requires a principal storage service for all ESXi hosts within the workload domain. Once a principal storage service has been provisioned and the workload domain has been created, supplemental storage can be provided to the workload domain using the vSphere Client.

### 3 About VI Workload Domains

In the VI Configuration wizard, you specify the storage, name, compute, and networking details for the VI workload domain. Based on the selected storage, you provide vSAN parameters, NFS share details, or VMFS on FC datastore information. You then select the hosts and licenses for the workload domain and start the creation workflow.

#### 4 Deploying NSX-T Edge Clusters

You can deploy NSX-T Edge clusters with 2-tier routing to provide north-south routing and network services in the management domain and VI workload domains.

#### 5 View Workload Domain Details

The Workload Domains page displays high level information about the workload domains in the Cloud Foundation system. CPU, memory, and storage utilized by the workload domain is also displayed here.

#### 6 Delete a VI Workload Domain

When you delete a workload domain, the clusters within that workload domain are deleted and the hosts are returned to the free pool.

#### 7 View vSphere Cluster Details

The cluster 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.

#### 8 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.

#### 9 Expand a Workload Domain

You can expand the management domain or a VI workload domain to add resources to support additional workloads or availability.

## Adding Virtual Machines to the Management Domain

You can add virtual machines to the management domain as desired. Commonly, these virtual machines provide infrastructure services such as backup or security throughout the solution. To prevent resource conflicts between the core Cloud Foundation services, these additional virtual machines are added to the Compute-ResourcePool. This resource pool is automatically created during bring-up for this purpose.

---

**Note** You must be careful when adding virtual machines to the management domain. You do not want to consume excessive resources that would obstruct standard operations. Excess capacity consumption can cause failures of virtual machine fail overs in the event of a host failure or maintenance action.

---

You can add capacity to the management domain by adding a host(s) in order to expand the management workload domain. To expand the management domain, see [Expand a Workload Domain](#).

#### Procedure

- 1 On the SDDC Manager Dashboard, navigate to **Inventory > Workload Domains**.
- 2 In the workload domains table, click **MGMT**.

- 3 On the MGMT page, click the **Services** tab.

- 4 Click the vCenter link.

This opens the vSphere Web Client for the management domain.

- 5 Create a VM.

See *Create a New Virtual Machine* in *vSphere Resource Management*.

---

**Note** Do not move any of the Cloud Foundation management VMs into the resource pool.

---

- 6 Move the VM to the resource pool.

See *Add a Virtual Machine to a Resource Pool* in *vSphere Resource Management*.

---

**Note** Do not move any of the Cloud Foundation management VMs to the newly created resource pool.

---

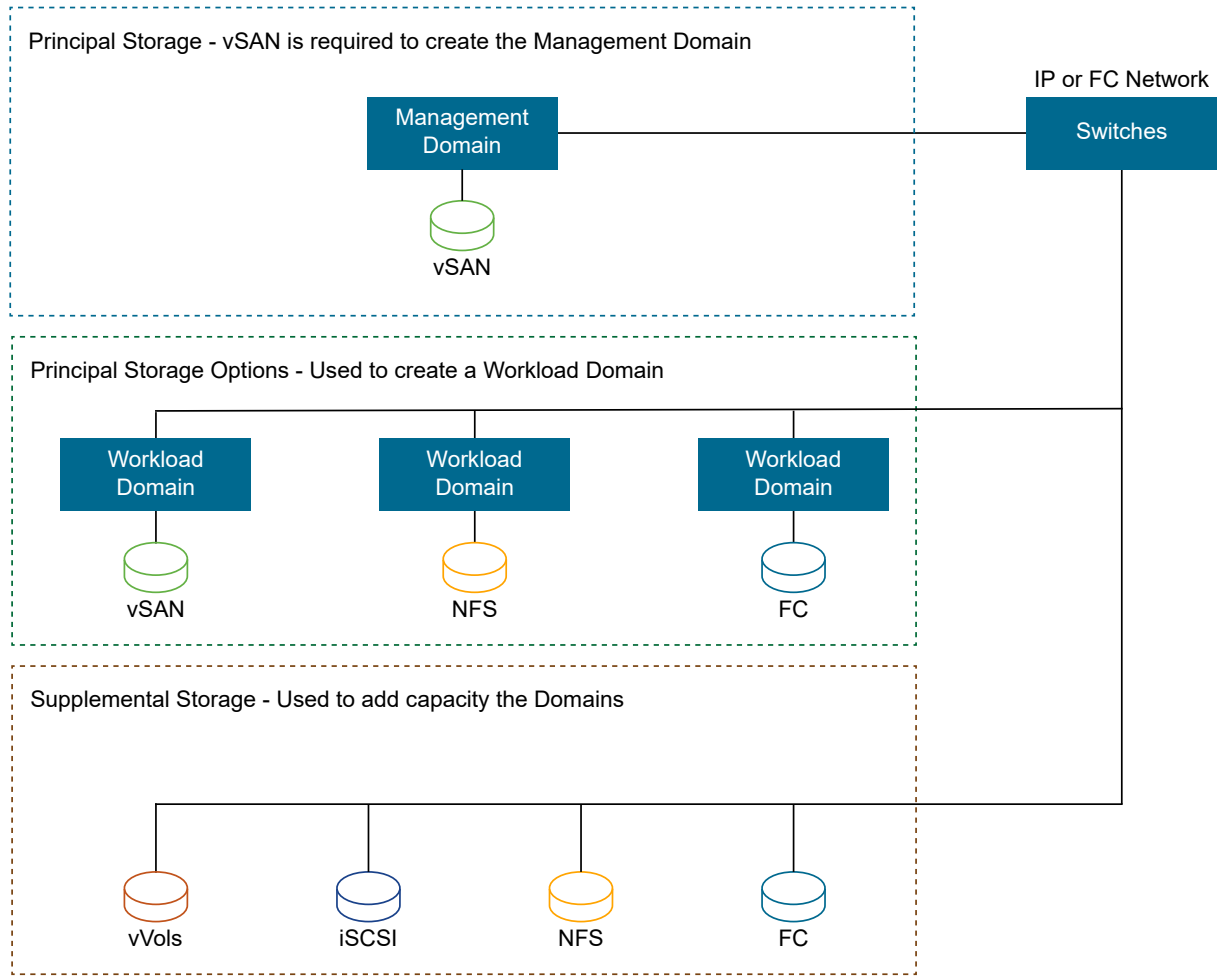
## Storage for the Management Domain and VI Workload Domains

To successfully create and manage a workload domain, Cloud Foundation requires a principal storage service for all ESXi hosts within the workload domain. Once a principal storage service has been provisioned and the workload domain has been created, supplemental storage can be provided to the workload domain using the vSphere Client.

Cloud Foundation uses and is validated against vSAN, NFS v3, and VMFS on FC for principal storage. The management domain requires vSAN for principal storage. You can use vSAN, NFS v3, or VMFS on FC for principal storage with VI workload domains. The type of principal storage used by the initial vSphere cluster in a VI workload domain is defined when the VI workload domain is created. For VI workload domains, you can add vSphere clusters that use a different type of principal storage than the type selected for the initial cluster when the VI workload domain was created. After a vSphere cluster has been deployed you cannot change to another storage type.

Using vSAN as the principal storage for all VI workload domains leverages the benefits of managing and maintaining a full software-defined stack. vSAN is also updated and patched through SDDC Manager lifecycle management. Updating and patching non-vSAN storage is a manual task and falls outside of SDDC Manager lifecycle management. To ensure supportability, non-vSAN storage and HBAs must be validated against the vSphere HCL.

Supplemental storage can be added to the management domain or a VI workload domain for the purposes of workload migration, backup and archive purposes. Cloud Foundation supports vVols, iSCSI, NFS (v3 or v4.1), and VMFS on FC as supplemental storage. See [vSphere Storage](#) for information about adding these storage types.



Both the principal and supplemental storage must be listed in the VMware Compatibility Guide for vSphere Storage and validated by the storage vendor as a VMware Cloud Foundation supplemental storage solution. Supplemental storage is not managed by or integrated into SDDC Manager processes. Because Cloud Foundation manages lifecycle processes automatically and without direct intervention, any storage provided to a workload domain must not modify the workload domain hosts in such a way that maintenance mode or graceful guest shutdown requests are prevented or dependent upon another separate process or action. Additionally, since workload domains can be scaled up or down by adding or removing ESXi hosts, any storage provided to a workload domain must not prevent the processes used by SDDC Manager to add or remove hosts in vSphere clusters.

Storage requiring the use of VMware Installation Bundles (VIBs) is supported, provided the storage presented is vSAN, NFS, or VMFS on FC. However, VIBs which must be installed or upgraded after deploying or upgrading ESXi are not managed or maintained by Cloud Foundation. Work with your storage vendor to manage these VIBs.

## About VI Workload Domains

In the VI Configuration wizard, you specify the storage, name, compute, and networking details for the VI workload domain. Based on the selected storage, you provide vSAN parameters, NFS share details, or VMFS on FC datastore information. You then select the hosts and licenses for the workload domain and start the creation workflow.

The workflow automatically:

- Deploys an additional vCenter Server Appliance for the new workload domain within the management domain.

By using a separate vCenter Server instance per workload domain, software updates can be applied without impacting other workload domains. It also allows for each 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 workload domain.
- Configures networking on each host.
- Configures vSAN, NFS, or VMFS on FC 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. In order to share an NSX Manager cluster, the workload domains must use the same update manager. The workload domains must both use vSphere Lifecycle Manager (vLCM) or they must both use vSphere Update Manager (VUM).
- During bring-up with application virtual networks (AVNs), Cloud Foundation creates a two-node NSX Edge cluster on the management domain for use by the vRealize Suite components. You can add additional NSX Edge clusters on the management domain. By default, workload domains do not include any NSX Edge clusters and are isolated. Add one or more Edge clusters to a workload domain to provide north-south routing and network services. See [Deploying NSX-T Edge Clusters](#).

---

**Note** Multiple Edge clusters cannot reside on the same vSphere cluster.

---

- NSX Managers deployed as part of a VI workload domain are configured to periodically get backed up to an SFTP server. By default, these backups are written to an SFTP server built into SDDC Manager, but you can register an external SFTP server for better protection against failures. See [Configure an External SFTP Server for File-Based Backups](#). SDDC Manager uses either the built-in or external SFTP server with all currently deployed NSX Managers and when deploying additional NSX Managers.
- Licenses and integrates the deployed components with the appropriate pieces in the Cloud Foundation software stack.

The result is a workload-ready SDDC environment.

---

**Note** You can only perform one workload domain operation at a time. For example, while creating a new workload domain, you cannot add a cluster to any other workload domain.

---

## Procedure

### 1 Prerequisites for a Workload Domain

This section lists pre-requisites for a VI workload domain.

### 2 Start the VI Configuration Wizard

Start the VI Configuration wizard and select the storage type for the workload domain.

## Prerequisites for a Workload Domain

This section lists pre-requisites for a VI workload domain.

- A DHCP server must be configured on the NSX-T host overlay (Host TEP) VLAN of the workload domain. When NSX-T creates Edge Tunnel End Points (TEPs) for the VI workload domain, they are assigned IP addresses from the DHCP server.
- A minimum of three hosts marked with the appropriate storage must be available in your Cloud Foundation inventory.
  - 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 an 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.

For information on adding hosts to your inventory, see [Chapter 7 Host Management](#).

- There must be a free pNIC on each host to be used for the 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 [Chapter 8 Cluster Image Management](#).
- If the management domain in your environment has been upgraded to a version different from the original installed version, you must download a VI workload domain install bundle for the current version before you can create a VI workload domain. See [Chapter 12 Download an Install Bundle](#).

- Decide on a name for your VI workload domain. Each VI workload domain must have a unique name. It is good practice to include 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-T Manager admin password

Although the individual Cloud Foundation components support different password requirements, VMware recommends that you set passwords following a common set of requirements across all components:

- Minimum length: 12
- Maximum length: 16
- At least 1 lowercase letter, 1 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
- Gather the information that you need for the workload domain creation workflow.

**Table 9-1. Information Required**

vCenter IP address and FQDN
Three NSX Managers IP addresses and FQDNs
NSX Manager Virtual IP (VIP) address and FQDN

- The IP addresses and Fully Qualified Domain Names (FQDNs) for the vCenter and NSX Manager instances must be resolvable by DNS.



- If you are using NFS storage for the 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 Cloud Foundation network. You must have read/write permission to the NFS share.

- If you are using VMFS on FC storage for the workload domain, you must create the datastore on the hosts and configure zoning and mounting of associated volumes.
- You must have valid license keys for the following products:
  - NSX-T Data Center
  - 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 workload domain. See [Chapter 5 License Management](#).

## Start the VI Configuration Wizard

Start the VI Configuration wizard and select the storage type for the workload domain.

### Procedure

- 1 On the SDDC Manager Dashboard, click **+ Workload Domain** and then click **VI Virtual Infrastructure**.
- 2 Select the storage type and click **Begin**.

## Specify Names and Choose an Update Manager

Provide names for the VI workload domain and organization and select the update manager for the workload domain. A VI workload domain can use vSphere Update Manager (VUM) or vSphere Lifecycle Manager (vLCM) as its update manager.

### Prerequisites

Verify that you have met the prerequisites described in [About VI Workload Domains](#).

### Procedure

- 1 Type a name for the VI workload domain, such as **sfo01**. The name must contain between 3 and 20 characters.

It is good practice to include location information in the name since resource object names (such as host and vCenter names) are generated based on the VI workload domain name.

- 2 (Optional) Type a name for the organization that requested or will use the virtual infrastructure, such as **Finance**. The name must contain between 3 and 20 characters.
- 3 Choose the update manager for the workload domain.

The update manager that you select for the workload domain cannot be changed later.

Option	Description
<b>Enable VUM (vSphere Update Manager)</b>	This is the default option. If you plan to use the workload domain for Kubernetes- Workload Management, VUM is the required update manager. See <a href="#">Chapter 10 Working with Workload Management</a> . If you enable VUM, you cannot use cluster images for firmware upgrades.
<b>Enable vLCM (vSphere Lifecycle Manager)</b>	Enable vLCM to use cluster images for firmware upgrades. A cluster image represents a desired software specification to be applied to all hosts in a cluster. See <a href="#">Chapter 8 Cluster Image Management</a> . vLCM-enabled workload domains do not support Kubernetes- Workload Management.

- 4 Click **Next**.

## Specify vSphere Cluster Details

Provide a name for the workload domain vSphere cluster. If you selected vSphere Lifecycle Manager, select a cluster image to apply to the hosts.

### Prerequisites

You must have a cluster image available if the workload domain is using vSphere Lifecycle Manager. See [Chapter 8 Cluster Image Management](#).

### Procedure

- 1 Enter a vSphere cluster name.
- 2 Select a cluster image from the drop-down menu.

The option is only available for workload domains that use vSphere Lifecycle Manager.

## Specify Compute Details

Specify the details for the vCenter Server that gets deployed for the workload domain.

### Procedure

- 1 Enter an IP address and FQDN for the workload domain's vCenter Server.
- 2 Enter and confirm a vCenter Server root password.
- 3 Click **Next**.

## Specify Networking Details

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.

Do not share an NSX Manager cluster between workload domains catering to different use cases that would require different NSX-T Edge cluster specifications and configurations.

See [VMware Configuration Maximums](#) for information about the maximum number of workload domains that can be managed by a single NSX Manager instance.

#### Procedure

- 1 On the Networking page of the wizard, choose to create a new NSX Manager cluster or reuse an existing one.

For the first VI workload domain, you must create an NSX Manager cluster.

- 2 If you are reusing an existing NSX Manager cluster, select the cluster.

The networking information for the selected cluster will display and cannot be edited. Skip to step 5.

---

**Note** In order to share an NSX Manager cluster, the workload domains must use the same update manager. The workload domains must both use vSphere Lifecycle Manager (vLCM) or they must both use vSphere Update Manager (VUM).

---

- 3 If you are creating a new NSX Manager cluster, enter the VLAN ID for the NSX-T host overlay (host TEP) network.
- 4 Provide the NSX Manager cluster details:
  - NSX Manager Virtual IP (VIP) address and FQDN
  - IP addresses and FQDNs for three NSX Managers (nodes)
  - NSX Manager Admin password
- 5 Click **Next**.

#### Select the vSAN Parameters

At the Storage step of the creation wizard, specify the availability you want provisioned for the VI workload domain. This page appears only if you are using vSAN storage for this workload domain.

Based on your selections, SDDC Manager will determine:

- The minimum number of hosts that it needs to fulfill those selections
- The specific hosts in your environment that are available and appropriate to fulfill those selections
- The virtual infrastructure features and their specific configurations that are needed to fulfill those selections

---

**Note** You can modify the vSAN configuration in vSphere without negatively affecting the Cloud Foundation configuration.

---

## Procedure

- 1 Specify the level of availability you want configured for this virtual environment.

The availability level determines the level of redundancy that is set for the assigned resources. For more information, see *Managing Fault Domains in Virtual SAN Clusters* in *Administering VMware Virtual SAN*.

Option	Description
0	<p>With this choice, the following vSAN parameters are used:</p> <ul style="list-style-type: none"> <li>■ Number of failures to tolerate: zero (0).</li> </ul> <p>Because vSAN requires a minimum of three hosts by default, three hosts are assigned to the virtual infrastructure.</p>
1	<p>With this choice, the following vSAN parameters are used:</p> <ul style="list-style-type: none"> <li>■ Number of failures to tolerate: one (1).</li> </ul> <p>Because vSAN requires a minimum of four hosts by default, four hosts are assigned to the virtual infrastructure. This is the default value.</p>
2	<p>With this choice, the following vSAN parameters are used:</p> <ul style="list-style-type: none"> <li>■ Number of failures to tolerate: two (2).</li> </ul> <p>Because vSAN requires a minimum of five hosts by default for this setting, five hosts are assigned to the virtual infrastructure.</p>

- 2 Click **Next**.

## Specify the VMFS on FC Datastore

If you are using VMFS on FC storage for the workload domain, you must specify the VMFS on FC datastore name.

## Procedure

- 1 On the Storage page, enter the name of the VMFS on FC datastore.
- 2 Click **Next**.

## Select Hosts

The Host Selection page displays available hosts along with hosts details. Hosts that are powered off, cannot be accessed via SSH, or have not been properly commissioned are not displayed.

- Select only healthy hosts.

To check a host's health, SSH in to the SDDC Manager VM using the **vcf** administrative user account and type the following command:

```
sudo /opt/vmware/sddc-support/sos --health-check
```

When prompted, enter the **vcf** user password. For more information, see [Chapter 18 Supportability and Serviceability \(SoS\) Utility](#)

- For optimum performance, you must select hosts that are identical in terms of memory, CPU types, and disks.

If you select unbalanced hosts, the UI displays a warning message, but you can proceed with the workload domain creation.

- You cannot select hosts that are in a dirty state. A host is in a dirty state when it has been removed from a cluster in a workload domain.

To clean a dirty host, re-image the host.

- For NSX for vSphere workload domains, all hosts in a cluster must be associated with the same network pool. For NSX-T workload domains, you can use the Cloud Foundation API to select hosts from different network pools, as long as those network pools have the same VLAN ID and MTU settings.

#### Procedure

- 1 Select the hosts for creating the VI workload domain.

For a vSAN VI workload domain with 0 or 1 availability, a minimum of three hosts is required. For a VI workload domain with 2 availability, a minimum of five hosts is required. When you select hosts with sufficient storage to form a VI cluster, the **Next** button is enabled.

You can add hosts with multiple active pNICs to the workload domain using the Cloud Foundation API.

The total resources based on the selected hosts are displayed.

- 2 Click **Next**.

### Specify NFS Storage Details

If you are using NFS storage for this workload domain, you must provide the NFS share folder and IP address of the NFS server.

#### Procedure

- 1 On the NFS Storage page, enter a name for the NFS datastore name.
- 2 Enter the path to the NFS share.
- 3 Enter the IP address of the NFS server.

---

**Note** When creating additional datastores for an NFS share and server, use the same datastore name. If you use a different datastore name, vCenter overwrites the datastore name provided earlier.

---

- 4 Click **Next**.

### Select Licenses

On the Licenses page select an available license for NSX-T, vSAN, and vSphere.

## Prerequisites

You must have added valid license keys for the following products:

- VMware vSAN (if using vSAN as the storage option)

NFS does not require a license

- VMware NSX-T Data Center

- VMware vSphere

Since vSphere and vSAN licenses are per CPU, ensure that you have sufficient licenses for the ESXi hosts to be used for the workload domain.

For information on adding license keys, see [Add License Keys for the Software in Your Cloud Foundation System](#).

## Procedure

- 1 Select a license key for each of the components in the VI workload domain.
- 2 Click **Next**.

## View Object Names

The Object Names page displays the vSphere objects that will be generated for the VI workload domain. Object names are based on the VI workload domain name.

## Procedure

- 1 Review the syntax that will be used for the vSphere objects generated for this domain.
- 2 Click **Next**.

## Review Details and Start the Creation Workflow

At the Review step of the wizard, review the information about the workload domain and start the creation workflow. You can also print the information or download a printable version to print later. It can take up to two hours for the domain to be created.

The Review page displays information about the resources and their configurations that are deployed when the workflow creates and deploys the virtual infrastructure for this workload domain.

The hosts that will be added to the workload domain are listed along with information such as the network pool they belong to, memory, CPU, and so on.

## Procedure

- 1 Scroll down the page to review the information.

## 2 Click **Finish** to begin the creation process.

The Workload Domains page appears, and a notification is displayed letting you know that VI workload domain is being added. Click **View Task Status** to view the domain creation tasks and sub tasks.

If a task fails, you can fix the issue and rerun the task. If the workload domain creation fails, contact VMware Support.

### Results

When the VI workload domain is created, it is added to the workload domains table.

The default end-state for a workload domain includes a network fabric with the following configuration

- Two NSX-T transport zones (one for VLAN and one for overlay) for each host
- NIOC and uplink profiles
- Four segments (previously known as logical networks); one each for management, vSAN, vMotion, and host overlay (host TEP)
- Some empty vDS and port groups

Do not edit or delete these.

### What to do next

Deploy and configure an NSX-T Edge cluster. See [Deploying NSX-T Edge Clusters](#) .

## Deploying NSX-T Edge Clusters

You can deploy NSX-T Edge clusters with 2-tier routing to provide north-south routing and network services in the management domain and VI workload domains.

An NSX-T Edge cluster is a logical grouping of NSX-T Edge nodes. These NSX-T Edge nodes run on a vSphere cluster, and provide north-south routing and network services for the management and VI workloads. NSX-T Data Center supports a 2-tier routing model. In the top tier is the 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. In the bottom tier is the tier-1 logical router. Northbound, the tier-1 logical router connects to a tier-0 logical router. Southbound, it connects to one or more logical switches.

During bring-up with application virtual networks (AVNs), Cloud Foundation creates a two-node NSX-T Edge cluster on the management domain for use by the vRealize Suite components. You can add additional NSX-T Edge clusters on the management domain to scale out or if you need custom configured services.

By default, workload domains do not include any NSX-T Edge clusters and are isolated. Add one or more Edge clusters to a workload domain to provide routing and network services.

You can add multiple NSX-T Edge clusters to the management or workload domains for scalability and resiliency.

NSX-T Data Center supports a maximum of 16 Edge clusters per NSX Manager cluster and 8 Edge clusters per vSphere cluster.

The north-south routing and network services provided by an NSX-T Edge cluster created for a workload domain are shared with all other workload domains that use the same NSX Manager cluster.

## Create an NSX-T Edge Cluster

You can add an NSX-T Edge cluster with 2-tier routing to the management domain or a workload domain to provide north-south routing and network services.

SDDC Manager does not enforce rack failure resiliency for Edge clusters. Make sure that the number of Edge nodes that you add to an NSX-T Edge cluster, and the vSphere clusters to which you deploy the Edge nodes, enable the Edge cluster to continue to provide Edge routing services in case of rack failure.

After you create an NSX-T Edge cluster, SDDC Manager does not support expanding or shrinking it by adding or deleting Edge nodes. If you need to expand or shrink an NSX-T Edge cluster, contact VMware Support.

This procedure describes how to use SDDC Manager to create an NSX-T Edge cluster with NSX-T 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](#).

### Prerequisites

- Separate VLANs and subnets are available for NSX-T Host Overlay (Host TEP) VLAN and NSX-T Edge Overlay (Edge TEP) VLAN. A DHCP server must be configured on the NSX-T Host Overlay (Host TEP) VLAN. You cannot use DHCP for the NSX-T Edge Overlay (Edge TEP) VLAN.
- NSX-T Host Overlay (Host TEP) VLAN and NSX-T Edge Overlay (Edge TEP) 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-T Edge cluster's Tier-0 gateway.
- DNS entries for the NSX-T Edge nodes are populated in the customer-managed DNS server.
- The vSphere cluster hosting an NSX-T Edge cluster must include hosts with identical management, uplink, host TEP, and Edge TEP networks (L2 uniform).
- You cannot deploy an Edge cluster on a vSphere cluster that is stretched. You can stretch an L2 uniform vSphere cluster that hosts an Edge cluster.



- The management network and management network gateway for the Edge nodes must be reachable.
- In Cloud Foundation 4.0, Workload Management supports one Tier-0 gateway per transport zone. When creating an Edge cluster for Workload Management, ensure that its overlay transport zone does not have other Edge clusters (with Tier-0 gateways) connected to it. Starting from Cloud Foundation 4.0.1, this limitation has been removed.

#### Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
- 2 In the **Virtual Infrastructure (VI)** 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 information for the NSX-T Edge cluster and click **Next**.

Setting	Description
Edge Cluster Name	Enter a name for the Edge cluster.
MTU	Enter the MTU for the Edge cluster. The MTU can be 1600-9000.
ASN	Enter the BGP ASN for the Edge cluster.
Tier 0 Name	Enter a name for the tier-0 gateway.
Tier 1 name	Enter a name for the tier-1 gateway.
Edge Cluster Profile Type	Select <b>Default</b> or, if your environment requires specific Bidirectional Forwarding Detection (BFD) configuration, select <b>Custom</b> .
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 a password.
Edge Admin Password	Enter and confirm a password.
Edge Audit Password	Enter and confirm a password.

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	Select <b>Workload Management</b> to create an Edge cluster that complies with the requirements for running Workload Management. See <a href="#">Chapter 10 Working with Workload Management</a> . If you select this option, you cannot modify the Edge form factor or Tier-0 service high availability settings. Select <b>Custom</b> if you want to modify those settings.
Edge Form Factor	<p>The default setting is <b>Large</b>.</p> <ul style="list-style-type: none"> <li>■ 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.</li> <li>■ Medium: 8 GB memory, 4 vCPU, 200 GB disk space. The NSX Edge Medium appliance size is suitable for a typical production environment.</li> <li>■ Large: 32 GB memory, 8 vCPU, 200 GB disk space. The NSX Edge Large appliance size is suitable for environments with load balancing.</li> </ul> <p>Workload management requires <b>Large</b>.</p>
Tier-0 Service High Availability	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 <b>Active-Active</b> . Some services are only supported in <b>Active-Standby</b> : NAT, load balancing, stateful firewall, and VPN. If you select <b>Active-Standby</b> , use exactly two Edge nodes in the Edge cluster.
Tier-0 Routing Type	Select <b>Static</b> or <b>EBGP</b> to determine the route distribution mechanism for the tier-0 gateway. If you select <b>Static</b> , you must manually configure the required static routes in NSX Manager. If you select <b>EBGP</b> , Cloud Foundation configures eBGP settings to allow dynamic route distribution.

7 Enter the NSX-T Edge node details for the first node and click **Add Edge Node**.

Setting	Description
Edge Node Name (FQDN)	Enter the FQDN for the Edge node. Each node must have a unique FQDN.
Management IP (CIDR)	Enter the CIDR for the management network. Each node must have a unique management IP.
Management Gateway	Enter the IP address for the management network gateway.
Edge TEP 1 IP (CIDR)	Enter the CIDR for the first Edge TEP. Each node must have a unique Edge TEP 1 IP.

Setting	Description
Edge TEP 2 IP (CIDR)	Enter the CIDR for the second 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 Edge TEP gateway.
Edge TEP VLAN	Enter the Edge TEP VLAN ID.
Cluster	Select a vSphere cluster to host the Edge node.
Cluster Type	<p>Select <b>L2 Uniform</b> if all hosts in the vSphere cluster have identical management, uplink, host TEP, and Edge TEP networks.</p> <p>Select <b>L2 non-uniform and L3</b> if any of the hosts in the vSphere cluster have different networks.</p> <p><b>Important</b> Cloud Foundation does not support Edge cluster creation on <b>L2 non-uniform and L3</b> vSphere clusters.</p>
First Uplink VLAN	<p>Enter the VLAN ID for the first uplink.</p> <p>This is a link from the NSX-T Edge node to the first uplink network.</p>
First 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)
ASN Peer	Enter the ASN for the first uplink peer. (EBGP only)
BGP Peer Password	Enter and confirm the BGP password. (EBGP only). A BGP password is required.
Second Uplink VLAN	<p>Enter the VLAN ID for the second uplink.</p> <p>This is a link from the NSX-T Edge node to the second uplink network.</p>
Second 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). A BGP password is required.

- 8 Click **Add More Edge Nodes** and enter the Edge node details.

A minimum of two NSX-T Edge nodes is required. You can have up to 10 NSX-T Edge nodes per Edge cluster.

- 9 When you are done adding NSX-T Edge nodes, click **Next**.

- 10 Review the summary and click Next.

SDDC Manager validates the NSX-T Edge node information.

- 11 If validation fails, use the **Back** button to edit your settings and try again.

To edit or delete any of the Edge nodes, click the three vertical dots next to an Edge node in the table and select an option from the menu.

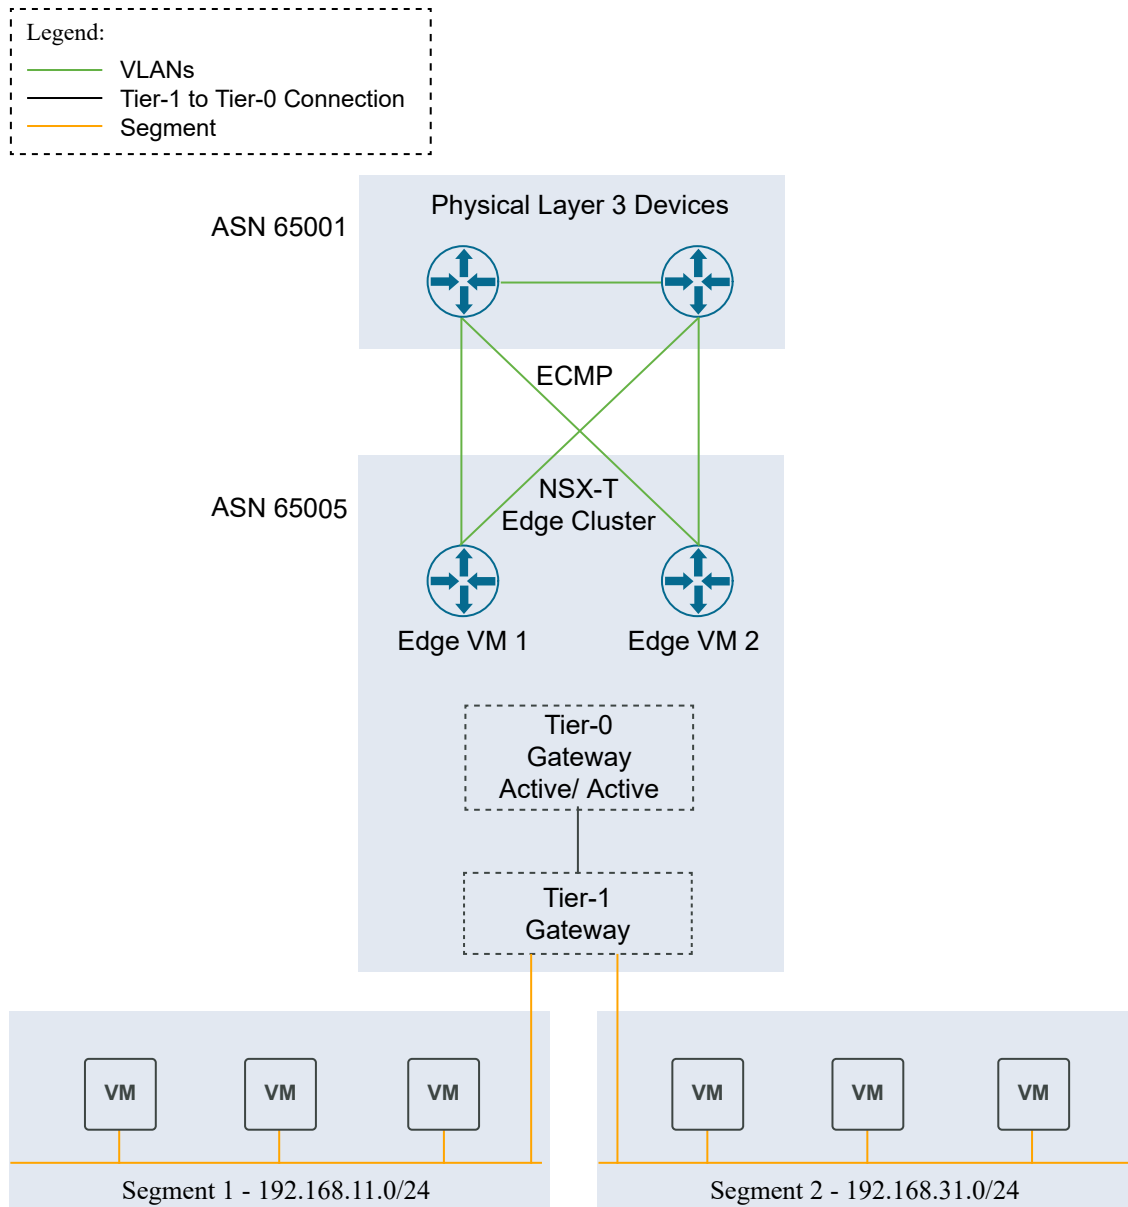
12 If validation succeeds, click **Finish** to create the NSX Edge cluster.

You can monitor progress in the Tasks panel.

### Example

The following example shows a scenario with sample data. You can use the example to guide you in creating NSX-T Edge clusters in your environment.

**Figure 9-1. Two-node NSX-T Edge cluster in a single rack**



### What to do next

In NSX Manager, you can create segments connected to the NSX-T Edge cluster's tier-1 gateway. You can connect workload VMs to these segments to provide north-south and east-west connectivity.

## View Workload Domain Details

The Workload Domains page displays high level information about the workload domains in the Cloud Foundation system. CPU, memory, and storage utilized by the workload domain is also displayed here.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.
- 2 In the workload domains table, click the name of the workload domain.

The domain details page displays CPU, memory, and storage allocated to the domain. The tabs on the page display additional information as described in the table below.

Tab	Information Displayed
Summary	Clusters in the workload domain and availability level for each cluster.
Services	<p>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 to reach the vSphere Web Client for that workload domain.</p> <p>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.</p>
Updates/Patches	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.
Security	Default certificates for the Cloud Foundation components. For more information, see <a href="#">Chapter 4 Certificate Management</a> .

### What to do next

You can add a cluster to the workload domain from this page.

## Delete a VI Workload Domain

When you delete a workload domain, the clusters within that workload domain are deleted and the hosts are returned to the free pool.

Deleting a VI workload domain also removes the components associated with the workload domain from the management domain. This includes the vCenter Server instance and the NSX Manager cluster.

---

**Note** If the NSX Manager cluster is shared with any other VI workload domains, it is not deleted.

---

The network pools used by the workload domain are not deleted as part of the workload domain deletion process and must be deleted separately.

---

**Caution** Deleting a workload domain is an irreversible operation. All clusters and VMs within the workload domain are deleted and the underlying datastores are destroyed.

---

It can take up to 20 minutes for a workload domain to be deleted. During this process, you cannot perform any operations on workload domains.

#### Prerequisites

- Back up the data on the workload domain. The datastores on the workload domain are destroyed when the workload domain is deleted.
- Migrate the VMs that you want to keep to another workload domain.
- Delete any workload VMs created outside Cloud Foundation before deleting the workload domain.
- Delete any NSX Edge clusters hosted on the workload domain. See [KB 78635](#).

#### Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 Hover your mouse in the workload domain row that where you want to delete.

When you select the workload domain, three vertical dots appear next to the name.

- 3 Click the dots and choose **Delete Domain**.

A confirmation window appears with details about the impact of deleting the workload domain, including how many hosts will be returned to the free pool.

- 4 Click **Delete Domain** to proceed.

The details page for the workload domain appears with a message indicating that the workload domain is being deleted. When the removal process is complete, the workload domain is removed from the domains table.

## View vSphere Cluster Details

The cluster 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.

**Procedure**

- 1 On the SDDC Manager Dashboard, 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 cluster detail page appears. The tabs on the page display additional information as described in the table below.

Tab	Information Displayed
Summary	Storage parameters configured on the vSphere cluster and organization name.
Hosts	Details about each host in the vSphere cluster. You can click a name in the FQDN column to access the host detail page.

**What to do next**

You can add or remove a host, or access the vSphere Client from this page.

## 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 the SDDC Manager Dashboard.

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.

**Prerequisites**

- Delete any workload VMs created outside Cloud Foundation or move them to another host.
- You cannot remove a host from a vSphere cluster if that vSphere cluster is hosting an Edge node.
  - Move the Edge node to another host in the same vSphere cluster. You cannot move an Edge node to a host that already hosts another Edge node.
  - If you cannot move the Edge node, delete the Edge node on the host. You cannot delete Edge nodes if doing so would result in an Edge cluster with fewer than two Edge nodes.
- You cannot remove a host from a vSphere cluster if that vSphere cluster was selected for Edge node placement and the Edge node deployment is still in pending state.

**Procedure**

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 In the workload domains table, click the name of the workload domain that you want to modify.

The detail page for the selected workload domain appears.

- 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 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.

**Results**

The host is removed from the workload domain and added to the free pool.

**What to do next**

Re-image the host so that 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. Datastores on the ESXi hosts in the deleted cluster are destroyed.

You cannot delete the last cluster in a workload domain. Instead, delete the workload domain. See [Delete a VI Workload Domain](#).

**Prerequisites**

- Delete any workload VMs created outside 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 vSphere cluster or shrink the NSX Edge cluster by deleting Edge nodes hosted on the vSphere 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](#).



### Procedure

- 1 On the SDDC Manager Dashboard, 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 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.

## 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 Cloud Foundation 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.

- Add a new vSphere cluster to a workload domain.

As workload domains support multiple vSphere clusters, you can add an additional cluster to an existing workload domain to provide for increased capacity and VM failover isolation.

## Add a Host to a vSphere Cluster in a Workload Domain

Adding an individual host to the management domain or a VI workload domain adds the resources of that host to the domain. You can add multiple hosts at a time to a domain.

### Prerequisites

- There must be a host available in the Cloud Foundation inventory. For information on adding a host to Cloud Foundation, see [Commission Hosts](#).
- Ensure that the host you want to add is in an active state.
- You must have a valid vSphere license specified in the Licensing tab of the SDDC Manager Dashboard with adequate sockets available for the host to be added. For more information, see [Add License Keys for the Software in Your Cloud Foundation System](#).

- 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 will be displayed. The warning will not prevent the expansion and can be dismissed if needed.
- The host you are adding must have the same type of principal storage as the existing hosts in the cluster. For the management domain, the host must use vSAN for principal storage. For VI workload domains, the host can use vSAN, NFS, or VMFS on FC for principal storage. A host using NFS for principal storage will automatically use the same NFS configuration as the other hosts in the cluster. For a host using VMFS on FC, you must create the datastore on the host and configure zoning and mounting of associated volumes before adding the host to a vSphere cluster.
- If the vSphere cluster hosts an NSX-T 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.

#### Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 In the workload domains table, click the name of the workload domain that you want to expand.

The detail page for the selected workload domain appears.

- 3 Click the **Clusters** tab.

- 4 Click the name of the vSphere cluster where you want to add a host.

- 5 Click **Actions > Add Host**.

The Add Hosts wizard appears.

- 6 Select the cluster expansion type.

This option only appears if the vSphere cluster hosts an NSX-T Edge cluster.

Option	Description
<b>L2 Uniform</b>	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.
<b>L2 non-uniform and L3</b>	Select if any of the hosts you are adding to the vSphere cluster have different networks than the existing hosts in the vSphere cluster.
	<b>Important</b> Cloud Foundation does not support adding hosts to <b>L2 non-uniform and L3</b> vSphere clusters that host an NSX-T Edge cluster.

## 7 Select the host you want to add to the vSphere cluster.

When you use the SDDC Manager UI to add a host, it must be associated with the same network pool as the other hosts in the cluster. The Cloud Foundation API supports adding hosts from different network pools to NSX-T workload domain clusters, as long as those network pools have the same VLAN ID and MTU settings.

For optimum performance, you should select hosts that are identical in terms of memory, CPU types, and disks to the other hosts in the vSphere cluster. If you select unbalanced hosts, the UI displays a warning message, but you can proceed with the workload domain creation.

You can add a host with multiple active pNICs to a multi pNIC-aware cluster using APIs.

## 8 Click **Next**.

## 9 Select the vSphere license you want to apply to the host.

## 10 Click **Next**.

## 11 Review the host and license details and click **Finish**.

The details page for the vSphere cluster appears with a message indicating that the host is being added. Wait until the action is complete before performing additional workload domain tasks.

# Add a vSphere Cluster to a Workload Domain

You can add a vSphere cluster to the management domain or to a VI workload domain through the Workload Domains page in the SDDC Manager Dashboard.

For the management domain, all vSphere clusters must use vSAN for principal storage. For VI workload domains, vSphere clusters can use vSAN, NFS, or VMFS on FC for principal storage. If a VI workload domain has multiple clusters, each clusters can use a different type of principal storage, as long as all hosts within a vSphere cluster use the same type.

### Prerequisites

- There must be at least three hosts available in the Cloud Foundation inventory. For information on adding a host to Cloud Foundation, see [Commission Hosts](#).
- Ensure that the hosts you want to add to the vSphere cluster are in an active state.
- You must have a valid vSphere and vSAN (if using vSAN storage) license specified in the Licensing tab of the SDDC Manager Dashboard with adequate sockets available for the host to be added. For more information, see [Add License Keys for the Software in Your Cloud Foundation System](#).
- A DHCP server must be configured on the NSX-T 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.
- From the UI, you can only select hosts on which pNICs 0 and 1 are active.

## Procedure

- 1 On the SDDC Manager Dashboard, click **Inventory > Workload Domains**.

The Workload Domains page displays information for all workload domains.

- 2 Use one of the following methods to get started.

- ◆ From the workload domain page:

- a Hover your mouse in the workload domain row that where you want to add a cluster.

A set of three dots appear on the left of the workload domain name.

- b Click these dots and then click **Add Cluster**.

The Add Cluster wizard appears.

- ◆ From the workload domain details page:

- a Click the name of the workload domain to go to the details page for that workload domain.

- b Click **Actions > Add Cluster**.

The Add Cluster wizard appears.

- 3 Select the storage type for the vSphere cluster and click **Begin**.

- 4 Enter a name for the vSphere cluster and click **Next**.

- 5 Select the cluster image to be applied to the vSphere cluster.

You can select a cluster image only if you are adding a vSphere cluster to a vLCM-enabled VI workload domain.

- 6 On the Networking page, enter the NSX-T host overlay (Host TEP) VLAN of the management domain and click **Next**.

- 7 If you selected vSAN storage for the vSphere cluster, the vSAN parameters page appears. Specify the level of availability you want configured for this cluster. The specified Failures To Tolerate (FTT) value determines the number of hosts required the cluster.

- 8 If you selected VMFS on FC storage for the vSphere cluster, enter the VMFS on FC datastore name.

- 9 Click **Next**.

- 10 On the Object Names page, review the syntax that will be used for the vSphere objects generated for this cluster and click **Next**.

- 11 On the Host Selection page, select hosts for the vSphere cluster.

When you use the SDDC Manager UI to add a cluster, all hosts must be associated with the same network pool. The Cloud Foundation API supports adding hosts from different network pools to NSX-T workload domain clusters, as long as those network pools have the same VLAN ID and MTU settings.

You can add hosts with multiple active pNICs to the vSphere cluster using APIs.

When you have selected the minimum number of hosts required for this cluster, the **Next** button is enabled.

**12** Click **Next**.

**13** If you selected NFS storage for the vSphere cluster, the NFS Storage page appears. Enter the datastore name, NFS share folder, and NFS server IP address.

**14** Click **Next**.

**15** On the Licenses page, select the vSphere and vSAN (if using vSAN storage) license to apply to this cluster.

**16** Click **Next**.

**17** On the Review page, review the vSphere cluster details and click **Finish**.

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.

# Working with Workload Management

# 10

With Workload Management, you can deploy and operate the compute, networking, and storage infrastructure for vSphere with Kubernetes. vSphere with Kubernetes transforms vSphere to a platform for running Kubernetes workloads natively on the hypervisor layer. When enabled on a vSphere cluster, vSphere with Kubernetes provides the capability to run Kubernetes workloads directly on ESXi hosts and to create upstream Kubernetes clusters within dedicated resource pools.

You begin the Workload Management deployment in Cloud Foundation, where the underlying infrastructure is validated. You then complete the deployment in the vSphere UI.

With Cloud Foundation 4.0.1, you can enable Workload Management on the management domain default cluster.

For more information on vSphere with Kubernetes, see [What Is vSphere with Kubernetes?](#).

This chapter includes the following topics:

- [Sizing Compute and Storage Resources for a vSphere with Kubernetes Workload Domain](#)
- [Deploy Workload Management](#)
- [View Workload Management Cluster Details](#)

## Sizing Compute and Storage Resources for a vSphere with Kubernetes Workload Domain

Compute and storage requirements for each component are key considerations when you size the solution.

Virtual Machine	Nodes	Total vCPUs	Total Memory	Total Storage
Supervisor Cluster control plane (small nodes - up to 2000 pods per Supervisor cluster)	3	12	48 GB	200 GB
Registry Service	N/A	7	7 GB	200 GB
Tanzu Kubernetes Cluster control plane (small nodes)	3 (per cluster)	6	12 GB	48 GB

Virtual Machine	Nodes	Total vCPUs	Total Memory	Total Storage
Tanzu Kubernetes Cluster worker nodes (small nodes)	3 (per cluster)	6	12 GB	48 GB
VMware NSX-T Edge node	2	16	64 GB	400 GB

## Deploy Workload Management

With Workload Management, you validate the underlying infrastructure for vSphere with Kubernetes and then complete the deployment in vSphere.

### Prerequisites

- A Workload Management ready NSX-T VI workload domain must have been deployed  
For a Workload Management ready NSX-T VI workload domain, you must have selected the **Enable VUM (vSphere Update Manager)** option on the Name page of the VI Configuration wizard. For more information, see [Specify Names and Choose an Update Manager](#).
- An NSX-T Edge cluster must have been deployed on the workload domain.  
Workload Management must have been selected on the Use Case page of the Add Edge Cluster wizard. See step 6 in [Create an NSX-T Edge Cluster](#).
- In Cloud Foundation 4.0, Workload Management supports one Edge cluster per transport zone, so ensure that the overlay transport zone connected to this cluster does not have other Edge clusters connected to it. Starting from Cloud Foundation 4.0.1, this limitation has been removed
- In a Cloud Foundation 4.0 environment, if the routing type for the Edge cluster is eBGP, you must manually configure the NCP route on Tier-0 router. See [Configure NSX Route Maps on Edge TO Router](#). This is not required for Cloud Foundation 4.0.1 and later releases.
- All hosts in the cluster where you want to deploy Workload Management must have a vSphere with Kubernetes license applied on them.
- The following IP addresses must have been defined. You must provide them when you complete Workload Management deployment in the vCenter UI.
  - Non-routable subnet for pod networking
  - Non-routable subnet for service IP addresses
  - Routable subnet for ingress
  - Routable subnet for egress

### Procedure

- 1 From the navigation bar, select **Solutions**.

**2** In the Kubernetes section, click **Deploy**.

The Workload Domains drop-down menu displays all Workload Management ready workload domains. Starting from Cloud Foundation 4.0.1, the management domain is also displayed for a consolidated architecture installation.

**3** Select the workload domain associated with the cluster where you want to deploy Workload Management.

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.

**4** From the list of compatible clusters on the workload domain, select the cluster where you want to deploy Workload Management.**5** Click **Next**.

On the Validation page, the following validations are performed.

- vCenter details (vCenter connectivity, objects, and version)
- Network validation (NSX-T details and version)
- Workload Management cluster compatibility

**6** Click **Next**.

The Review page displays information about the VI workload domain where you are deploying Workload Management. The remaining deployment is to be completed on the vSphere UI.

Note down the values for the following fields as you have to enter these values in the vSphere UI. These fields are marked with a check box on the UI.

- Cluster where you are deploying Workload Management
- DNS server for the cluster
- NTP server for the cluster
- Edge cluster on the workload domain where you are deploying Workload Management
- vSphere Distributed Switch for the workload domain

**7** Click **Complete in vSphere**.

The Workload Management page in the vSphere UI is displayed. The vSphere Center for the VI workload domain where you are deploying Workload Management is selected by default. Continue the deployment here. See [Create and Configure a Supervisor Namespace](#).

## Configure NSX Route Maps on Edge T0 Router

When you deploy Workload Management in Cloud Foundation 4.0, the route maps created on the NSX-T Edge Tier-0 router in eBGP mode contains an IP prefix with only a deny rule. This blocks



routes from getting advertised to the ToR switches. This issue does not exist in Cloud Foundation 4.0.1 and later releases.

If you are using the Edge cluster only for Kubernetes - Workload Management, follow option 1 and deactivate tier-1 route advertisements. If you are using the Edge cluster for additional tasks, follow option 2 and create a new allow rule.

### Option 1: Deactivate Advertisements of Tier-1 Connected Networks through Tier-0 Router

Networks connected to tier-1 router are not advertised from tier-0 gateway to outside networks.

- 1 Navigate to the NSX-T UI for the workload domain where you are deploying Workload Management.
- 2 In the Networking section, click **Tier-0 Gateway**.
- 3 Click the horizontal ellipsis and then click **Edit**.
- 4 Expand Route Re-Distribution and click the horizontal ellipsis to the right of Route Re-Distribution.
- 5 Click the horizontal ellipsis to the left of default and click **Edit**.
- 6 Click L2 in the Route Re-distribution column.
- 7 In the Advertised Tier-1 Subnets section, click the **Connected interfaces and Segments** checkbox (so that it is not selected).
- 8 Click **Apply** and then click **Save**.

### Option 2: Create New Allow Rule and Apply it to Route Re-redistribution

When you deploy Workload Management, a new deny rule is appended to the route map. So you must add a new permit rule to the route map.

allow any IP prefix list and route map and apply it to the route redistribution rule as the last rule.

- 1 Navigate to the NSX-T UI for the workload domain where you are deploying Workload Management.
- 2 In the Networking section, click the horizontal ellipsis to edit **Tier-0 Gateway**.
- 3 Create a new IP prefix list.
  - a Expand Routing.
  - b Click 1 next to IP Prefix Lists.
  - c In the Set IP Prefix List dialog box, click **Add IP Prefix List**.
  - d Type a name (for example, test) and click **Set**.
  - e Click **Add Prefix**.
  - f In Network, click **Any** and in Action, select **Permit**.

- g Click **Apply** and then click **Save**.
- 4 Create a route map for the IP prefix list created in step 3.
  - a Click **Set** next to Route Map.
  - b Click **Add Route Map**.
  - c Add new match criteria with IP prefix.
  - d Select the IP prefix created in step 3 and action **Permit**.
  - e Click **Apply** and then click **Save**.
- 5 Apply edited route map to route re-distribution.
  - a On the **Tier-0 Gateways** page, expand **Route Re-Distribution** and click the horizontal ellipsis next to Router Re-distribution.
  - b Click the horizontal ellipsis to the left of default and click **Edit**.
  - c From the drop-down in the Route Map column, select the route map you created in step 4.
  - d Click **Apply** and then click **Save**.
  - e Click **Close Editing**.

## View Workload Management Cluster Details

The Workload Management page displays clusters where Workload Management has been deployed. The status of each cluster, number of hosts in the cluster, and associated workload domain is also displayed.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Solutions**.
- 2 In the Kubernetes section, click **View Details**.
- 3 Click vSphere Workload Management Clusters to see cluster details in vSphere.

# Deploy vRealize Suite Lifecycle Manager in Cloud Foundation

# 11

Before you can deploy vRealize Log Insight, vRealize Operations, or vRealize Automation, you must deploy vRealize Suite Lifecycle Manager.

By default, VMware Cloud Foundation uses NSX-T Data Center to create application virtual networks (AVNs) and deploys vRealize Suite Lifecycle Manager to these AVNs. If you deactivate AVNs during bring-up, you can deploy vRealize Suite Lifecycle Manager to a VLAN-backed network as described in [KB 78608](#).

Cloud Foundation does not automate the deployment and lifecycle management of the other vRealize Suite components. You can use vRealize Suite Lifecycle Manager to deploy and manage those components.

## Prerequisites

- Download the vRealize Suite Lifecycle Manager installation package from the VMware Depot to the local bundle repository. See [Chapter 12 Download an Install Bundle](#).
- Allocate an IP address for the vRealize Suite Lifecycle Manager virtual appliance and prepare forward/reverse DNS records.

## Procedure

- 1 On the **SDDC Manager Dashboard**, navigate to **Administration > vRealize Suite**.
- 2 Click **vRealize Suite Lifecycle Manager** and click **Deploy**.

The **vRealize Lifecycle Manager Installation Prerequisites** page displays the prerequisites that you must complete before beginning the installation.

- 3 Review the readiness of each prerequisite and verify by selecting each adjacent check box.  
When all the boxes are selected, the **Begin** button is activated.

- 4 Click **Begin**.

- 5 On the **Network Settings** page, review the settings and click **Next** to continue.

- 6 On the **Virtual Appliance Settings** page, enter the settings and **Next** to continue.

Setting	Description
FQDN	Enter the FQDN for the vRealize Suite Lifecycle Manager virtual appliance.
System Administrator	Create and confirm a password for the vRealize Suite Lifecycle Manager system administrator (for example, admin@localuser). This is the credential that allows SDDC Manager to connect to the vRealize Suite Lifecycle Manager system.
SSH Root Account	Create and confirm a password for the vRealize Suite Lifecycle Manager virtual appliance root account.

- 7 On the **Review Summary** page, review the installation configuration settings.
- 8 Click **Finish**.

SDDC Manager validates the inputs and reports any errors or warnings.

---

**Note** If necessary, you can use the **Back** button to return to preceding pages and modify settings.

---

- 9 Address any validation issues and then click **Finish**.

The **vRealize Suite Lifecycle Manager** page displays with the following message: `Deployment in progress`. If the deployment fails, this page displays a deployment status of `Failed`. In this case, you can **Restart Task** or **Uninstall**.

- 10 (Optional) Click **View Status in Tasks** to view the details of the deployment in progress or a deployment failure.

The **Tasks** panel opens at the bottom page. You can open individual tasks to view details.

- 11 (Optional) After the successful deployment of vRealize Suite Lifecycle Manager, click the vRealize Suite Lifecycle Manager link below the page title.

The vRealize Suite Lifecycle Manager user interface opens in a new browser tab.

#### What to do next

You can now deploy vRealize Log Insight, vRealize Operations, and vRealize Automation. See [Deployment for Cloud Operations and Automation in the First Region](#).

# Download an Install Bundle

# 12

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 Cloud Foundation installation. It includes software bits for vCenter Server and NSX-T Data Center.
- The vRealize Suite Lifecycle Manager install bundle is used for deploying vRealize Suite Lifecycle Manager.

This section describes the procedure for downloading install bundles from SDDC Manager. To download install bundles using a proxy server or in an offline mode, refer to the Download Bundles section in the *VMware Cloud Foundation Lifecycle Management*.

## Procedure

- 1 Log in to your My VMware Account.
  - a On the SDDC Manager Dashboard, click **Administration > Repository Settings**.
  - b Click **Authenticate**.
  - c Type your My VMware user name and password.
  - d Click **Authorize**.
- 2 On the SDDC Manager Dashboard, click **Repository > Bundles** in the left navigation pane. All available bundles are displayed. Install bundles display an Install Only Bundle label.
- 3 For the bundle you want to download, do one of the following:
  - Click **Download Now**.

The bundle download begins right away.
  - Click **Schedule Download**.

Select the date and time for the bundle download and click **Schedule**.
- 4 Navigate to **Repository > Download History > to see the downloaded bundles**.

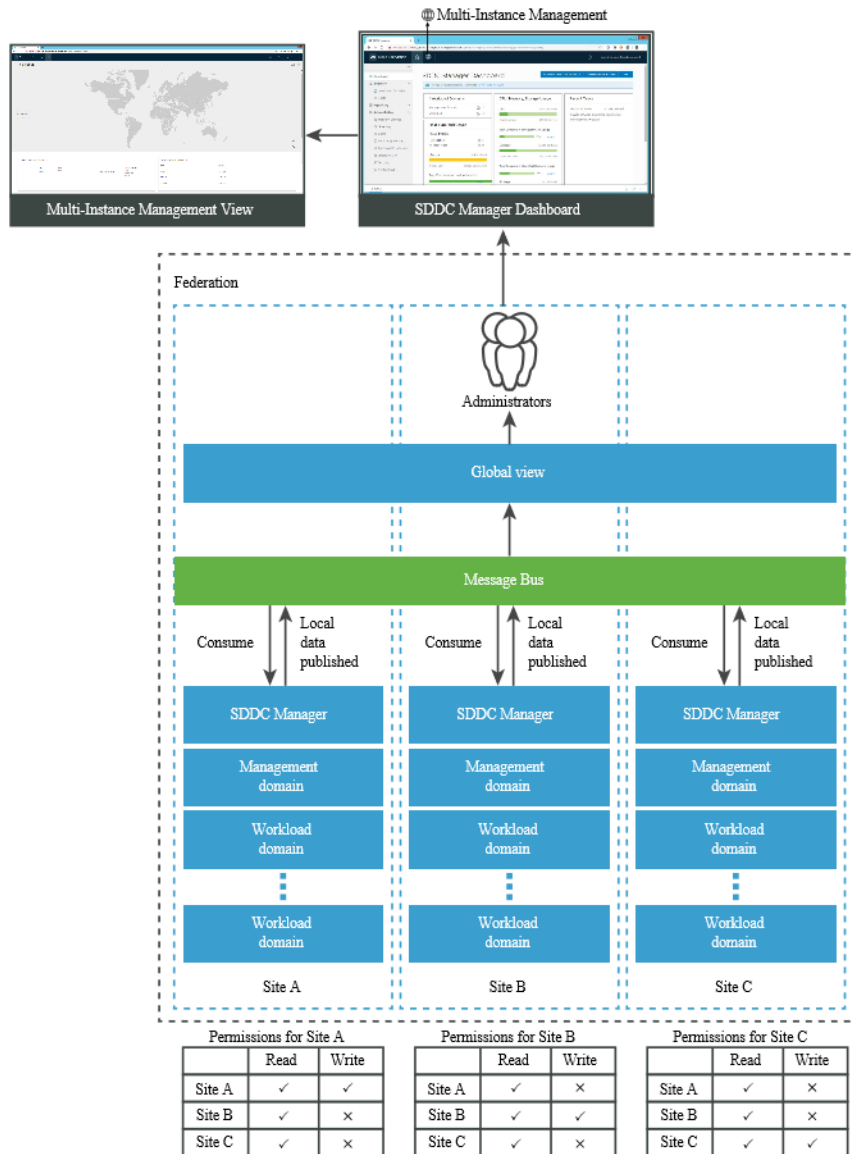
# Multi-Instance Management

# 13

With the Multi-Instance Management feature, you can monitor and manage multiple Cloud Foundation instances from a single console.

Multiple Cloud Foundation instances can be managed together by grouping them into a federation, such that each member can view information about the entire federation and the individual instances within it. Federation members can view inventory across the Cloud Foundation instances in the federation as well as the available and used capacity (CPU, memory, and storage). This allows you to maintain control over the different sites and ensure that they are operating with the right degree of freedom and meeting compliance regulations for your industry. It also simplifies patch management by showing the number of patches available across sites in the global view.

Federation members communicate with each other via a message bus. Each participant publishes their local data to the message bus and the remaining participants can read this data for global visibility across the federation.



An instance can see details about the federation only if it is a member of the federation, and can belong only to a single federation at a time. It is possible to create multiple federations within an organization; however, there is no global visibility between federations. For example, it might be desirable to have a dev-test federation and a production federation. In such an example, members of dev-test can see other dev-test members but they are not able to see production members.

Federation members can either be controllers or regular members. A controller member has capabilities of a regular member and runs some additional message bus components to allow multi-instance management to work.

A controller member can invite other instances to become members as controller or regular members. The controller role can be granted to a maximum of three instances within a federation. High Availability of multi-instance management functionality is only possible when there are exactly three controllers in the federation. Though the controller members can be at any three sites in the federation, it is recommended that each controller is in a different availability zone.

The instance who created the federation is automatically granted the controller role. If you only have two instances in the federation, there is no need to create both as controllers. Multi-Instance management works with two Cloud Foundation sites; however, if one fails then the multi-instance capability is not available on the other site.


This chapter includes the following topics:

- [About the Multi-Instance Management Dashboard](#)
- [Create a Federation](#)
- [Invite a Cloud Foundation Instance to Join a Federation](#)
- [Join a Federation](#)
- [Leave a Federation](#)
- [Dismantle a Federation](#)

## About the Multi-Instance Management Dashboard

The Multi-Instance Management Dashboard displays the inventory and capacity across the federation.

You access the Multi-Instance Management Dashboard by clicking the Multi-Instance View icon in the top left corner of the SDDC Manager Dashboard.




Before a federation is created, the dashboard displays a create and join option.

### Welcome to Multi-Instance Management

Please select one of the following based on your role.


#### Create a Federation



Only 1 user should create a federation for a given organization. This instance will become the first Controller. ⓘ

[CREATE](#)

#### Join a Federation

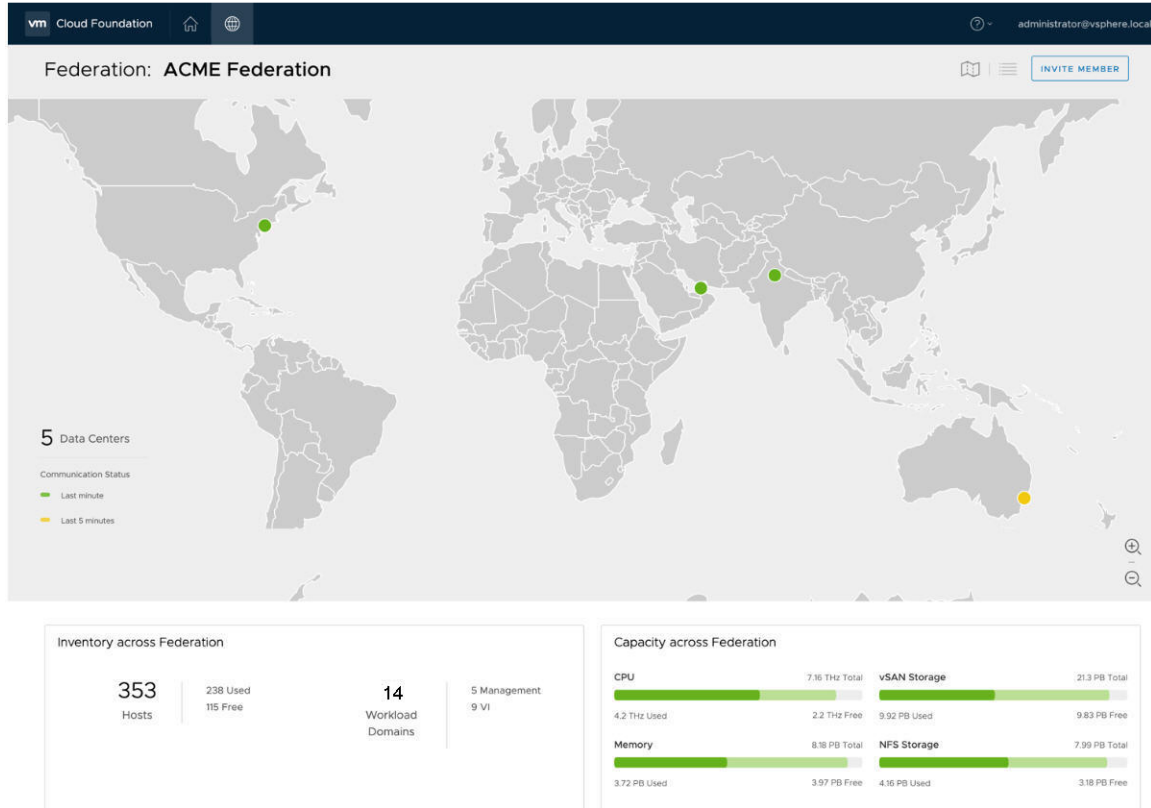


Most users will be joining an established federation by invitation. Please refer to the instructions you received in order to join.

[JOIN](#)



After a federation is created, the Multi-Instance Management Dashboard displays a world map showing the federation members as dots on the map.

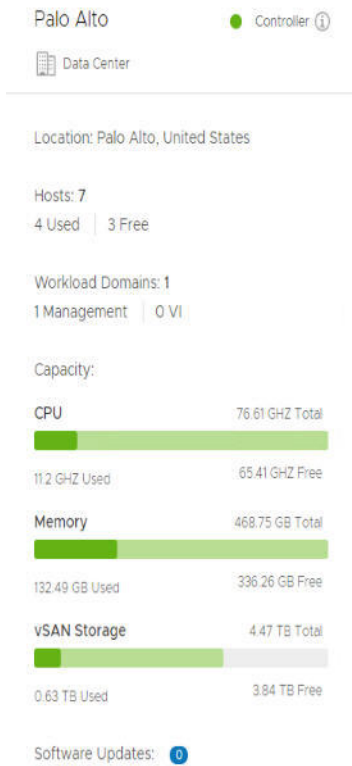


The dot color depends on the communication status between the federation members - green if they communicated within the last two minutes, yellow if they communicated within the previous five minutes, and red if they have not communicated for more than ten minutes. You can see the following information here:

- Hover over a dot to see the member name and location.



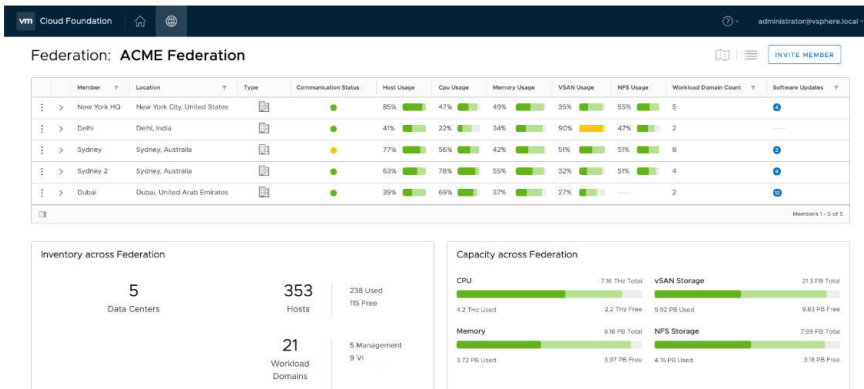
- Click the member location dot to open a panel on the right side with detailed information about the Cloud Foundation instance. The panel also displays available software updates.



The Inventory section in the bottom half of the dashboard displays the number of hosts and workload domains along with a breakdown of the workload domain type. The capacity section displays the used and available CPU, memory, and storage across the federation.




Click the Table icon at the top right of the dashboard to display member information in a grid format. Information for all federation members is displayed in a tabular format.



Clicking the arrow (➤) next to the member name displays the CPU, memory, and storage usage for that member.

### Federation: ACME Federation



Clicking MGMT takes you to the management domain of that member.

## Create a Federation

A federation is a group of Cloud Foundation instances, such that each member can view information about the other Cloud Foundation instances in the group. The federation creator is granted the controller role by default.

You can create multiple federations within your organization, but global visibility is available only within a federation. Members can belong to only a single federation at a time.

See [VMware Configuration Maximums](#) for information about the maximum number of SDDC Manager instances that can be managed using Multi-Instance Management

### Prerequisites

- The SDDC Manager certificate requires TLS Web server authentication and TLS Web client authentication extended key usage (EKU). If you are using Microsoft CA or a third-party CA, you must ensure that the Certificate Authority template is configured with both these EKUs.
- Retrieve the FQDN of your SDDC Manager.

### Procedure

- 1 On the SDDC Manager dashboard, click the Multi-Instance View (🌐) icon at the top of the window.

## 2 Click **Create**.

**Create Federation**

Establishing a federation requires two steps:

1. Register the first member (who is designated as a Controller).
2. Invite additional members to the federation.

**Responsibilities of the Controller:**

- Only Controllers may invite new members to a federation.
- Controllers can only dismantle a federation once all its members have left the federation.
- Providing high availability.

Federation Name	ACME Federation
Member Name ⓘ	Palo Alto
FQDN	sddc-manager.vrack.vsphere.local
Country	United States ▼
City	Palo Alto ▼

- 3 Enter a name for the federation.
- 4 Enter a display name for the member. You may want to base this on the location of this Cloud Foundation instance.
- 5 Type the FQDN of the SDDC Manager.
- 6 Select the city and country of this Cloud Foundation instance and click **Create**.

### Results

It can take a few minutes for the federation to be created. After the federation is created, the Multi Instance Management Dashboard is displayed. The federation location is marked with a green dot on the world map. You can zoom in or out of the map.

The dashboard also displays the inventory (hosts and workload domains) and capacity (CPU, memory, and storage) across the federation. These details are updated when additional members join the federation.

### What to do next

Invite a Cloud Foundation instance to join the federation.

## Invite a Cloud Foundation Instance to Join a Federation

You can invite Cloud Foundation instances to join a federation. They can be invited as a controller or a regular member. High Availability of multi-instance management functionality is only possible when there are exactly three controllers in the federation. Though the controller members can

be at any three sites in the federation, it is recommended that each controller is in a different availability zone.

### Prerequisites

You must be a controller in the federation and have the FQDN of the member you are inviting.

### Procedure

- 1 On the top right corner of the Multi-Instance Management dashboard, click **Invite Member**.
- 2 Enter the SDDC Manager FQDN of the member you are inviting and click **Check Certificate**.  
The invited member's certificate thumbprint is displayed.
- 3 Validate the thumbprint and click **Confirm fingerprint**.
- 4 Click **Next**.
- 5 Select the check box on the High Availability page if you want to designate the controller role to the member.
- 6 Click **Next**.

The Instructions page displays the URL that the invited member needs to access.

Invite Member

1 Enter member FQDN

2 High Availability

3 Member Instructions

High Availability ⓘ

✓

The federation already has a maximum of 3 controllers. High availability is in effect.

In order to ensure high availability performance, a federation must have exactly 3 controllers. Before deciding to designate a controller, please be aware of how many already exist for this federation.

☐ Designate this member as a controller. ⓘ

- 7 Click **Copy Info** to copy the information displayed on this page or copy the URL manually and send it to the member through an offline method.

### What to do next

The invitation and joining process is a coordinated effort between the invitee controller and joining member. An additional dot on your Multi-Instance Management Dashboard indicates that the member you invited is joining the federation. When a controller joins a federation, it can take a few minutes for the federation to stabilize.

## Join a Federation

You can join a federation as a controller or member depending on the assigned role in the invitation. An invitation is valid for ten days. You must request a new invitation after this period. If a new invitation is generated for the same site, only the latest invite is valid.

VMware, Inc.

101

## Prerequisites

- Verify that the federation is healthy - there should be no red dots on the world map of the Multi-Instance Management Dashboard. A red dot indicates that SDDC Manager is unable to communicate with that member. A controller must remove the member using the leave API. See the *VMware Cloud Foundation 4.0 API Reference Guide*.
- The SDDC Manager certificate requires TLS Web server authentication and TLS Web client authentication extended key usage (EKU). If you are using Microsoft CA or a third-party CA, you must ensure that the Certificate Authority template is configured with both these EKUs.
- Retrieve the FQDN of your SDDC Manager.

## Join a Federation by Clicking an Invitation

If you join a federation by clicking the invitation you received, federation details are pre-populated in the UI.

### Prerequisites

Retrieve the invitation you received.

### Procedure

- 1 Click the URL in the invitation you received.

The Join Federation window displays the role assigned in the invitation, the FQDN of the invited member, token, and the FQDN of the controller member who invited you to join the federation.

#### Join Federation

✔ Certificate is validated successfully. ✕

Member Name ①	<input type="text" value="Delhi"/>
Member Role	<input style="border-bottom: 1px solid #ccc;" type="text" value="Member"/>
FQDN ①	<input type="text" value="delhi.mydomain.local"/>
Country	<input style="border-bottom: 1px solid #ccc;" type="text" value="India"/>
City	<input style="border-bottom: 1px solid #ccc;" type="text" value="Delhi"/>
Token ①	<input type="text" value="jhdjfJHJGDJJKKDF57642j4JHBDkjndnk"/>
FQDN of Controller ①	<input type="text" value="newyork.mydomain.local"/>

- 2 Click **Join**.

## Results

The join process is initiated. If you see an error, resolve the issue and then request a new invitation and follow the steps described above. After the join process is successful, your Cloud Foundation instance becomes a member of the federation.

After a controller joins or leaves a federation, Kafka is restarted on all controllers in the federation. It can take several minutes for the federation to stabilize even after the dashboard is refreshed. If an operation performed on the dashboard during this time fails, re-try the operation.

## Join a Federation through the Multi-Instance Management Dashboard

You can join a federation through the Multi-Instance Management Dashboard.

### Procedure

- 1 Click **Join Federation** on the Multi-Instance Management Dashboard.
- 2 Type a display name for the site to be added.
- 3 Select the member role as indicated in the invitation you received.
- 4 Type the FQDN of your SDDC Manager.
- 5 Select the country and city for your site.
- 6 Type the token as indicated in the invitation you received.
- 7 Type the FQD of the controller who invited you.
- 8 Click **Join**.

## Results

The join process is initiated. If you see an error, resolve the issue and then request a new invitation and follow the steps described above. After the join process is successful, your Cloud Foundation instance becomes a member of the federation.

After a controller joins or leaves a federation, Kafka is restarted on all controllers in the federation. It can take several minutes for the federation to stabilize even after the dashboard is refreshed. If an operation performed on the dashboard during this time fails, re-try the operation.

## Leave a Federation

Leaving a federation removes the Multi-Instance Management view from your SDDC Manager Dashboard.

If you are a controller, you can leave a federation only if there is at least one more controller in the federation. If you are the only controller member in a federation, you must dismantle a federation instead of leaving it.

### Prerequisites

Verify that the federation is healthy - there should be no red dots on the world map of the Multi-Instance Management Dashboard. A red dot indicates that SSDC Manager is unable to communicate with that member. A controller must remove the member using the leave API. See the *VMware Cloud Foundation 4.0 API Reference Guide*.

### Procedure

- 1 Click the Grid icon at the top right of the Multi-Instance Management Dashboard.
- 2 In the member table, click the dot icon next to your member name and click **Leave Federation**.
- 3 Type the federation name and click **Leave**.

### What to do next

Do not perform any operation for a few minutes after leaving a federation.

## Dismantle a Federation

You can dismantle a federation if you are the last controller member in the federation. Only members with the controller role can dismantle a federation.

### Procedure

- 1 Click the Grid icon at the top right of the Multi-Instance Management Dashboard.
- 2 In the member table, click the dot icon next to your member name and click **Dismantle Federation**.
- 3 Type the federation name and click **Dismantle**.

### What to do next

After the federation is dismantled, the Create Federation screen is displayed instead of the Multi-Instance Management Dashboard.



# Stretching Clusters

# 14

You can stretch an NSX-T cluster in the management domain or in a VUM-based VI 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 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 primary availability zone goes down.

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.

This chapter includes the following topics:

- [About Availability Zones and Regions](#)
- [Stretched Cluster Requirements](#)
- [Deploy and Configure vSAN Witness Host](#)
- [Stretch a Cluster](#)
- [Expand a Stretched Cluster](#)
- [Unstretch a Cluster](#)
- [Replace a Failed Host in a Stretched Cluster](#)

## 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). 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 management, Uplink 01, Uplink 02, and Edge Overlay networks in each availability zone must be stretched to facilitate failover of the NSX-T Edge appliances between availability zones. The Layer 3 gateway for the management and Edge Overlay networks must be highly available across the availability zones.

**Table 14-1. Management Domain VLAN and IP Subnet Requirements**

Function	Availability Zone 1	Availability Zone 2	VLAN ID	IP Range	HA Layer 3 Gateway	Recommended MTU
Management (AZ1)	✓	✓	1611 (Stretched)	172.16.11.0/24	✓	1500
vSphere vMotion	✓	X	1612	172.16.12.0/24	✓	9000
vSAN	✓	X	1613	172.16.13.0/24	✓	9000
NSX-T Host Overlay	✓	X	1614	172.16.14.0/24	✓	9000
NSX-T Edge Uplink01	✓	✓	2711 (Stretched)	172.27.11.0/24	X	9000

**Table 14-1. Management Domain VLAN and IP Subnet Requirements (continued)**

Function	Availability Zone 1	Availability Zone 2	VLAN ID	IP Range	HA Layer 3 Gateway	Recommended MTU
NSX-T Edge Uplink02	✓	✓	2712 (Stretched)	172.27.12.0/24	X	9000
NSX-T Edge Overlay	✓	✓	2713 (Stretched)	172.27.13.0/24	✓	9000
Management (AZ2)	✓	✓	1621 (Stretched)	172.16.21.0/24	✓	1500
vSphere vMotion	X	✓	1622	172.16.22.0/24	✓	9000
vSAN	X	✓	1623	172.16.23.0/24	✓	9000
Host Overlay	X	✓	1624	172.16.24.0/24	✓	9000

**Note** If VLAN is stretched between AZ1 and AZ2, the Layer 3 network must also be stretched between the two AZs.

**Table 14-2. Workload Domain VLAN and IP Subnet Requirements**

Function	Availability Zone 1	Availability Zone 2	VLAN ID	IP Range	HA Layer 3 Gateway
Management (AZ1)	✓	X	1631	172.16.31.0/24	✓
vSphere vMotion	✓	X	1632	172.16.32.0/24	✓
vSAN	✓	X	1633	172.16.33.0/24	✓
Host Overlay	✓	X	1634	172.16.34.0/24	✓
Management (AZ2)	X	✓	2731	172.27.31.0/24	✓
vSphere vMotion	X	✓	2732	172.27.32.0/24	✓
vSAN	X	✓	2733	172.16.33.0/24	✓
Host Overlay	X	✓	1621	172.16.21.0/24	✓

**Note** If VLAN is stretched between AZ1 and AZ2, the Layer 3 network must also be stretched between the two AZs.

## Networking for Multiple Availability Zones

There are specific physical data center network requirements for a topology with multiple availability zones.

**Table 14-3. Physical Network Requirements for Multiple Availability Zone**

Component	Requirement
MTU	<ul style="list-style-type: none"> <li>■ VLANs which are stretched between availability zones must meet the same requirements as the VLANs for intra-zone connection including MTU.</li> <li>■ MTU value must be consistent end-to-end including components on the inter zone networking path.</li> <li>■ Set MTU for all VLANs and SVIs (management, vMotion, Geneve, and Storage) to jumbo frames for consistency purposes. Geneve overlay requires an MTU of 1600 or greater.</li> </ul>
Layer 3 gateway availability	For VLANs that are stretched between available zones, configure data center provided method, for example, VRRP or HSRP, to failover the Layer 3 gateway between availability zones.
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"> <li>■ For VLANs that are stretched between availability zones, traffic flows in and out of a single zone. Local egress is not supported.</li> <li>■ For VLANs that are not stretched between availability zones, traffic flows in and out of the zone where the VLAN is located.</li> <li>■ For NSX-T virtual network segments that are stretched between regions, traffic flows in and out of a single availability zone. Local egress is not supported.</li> </ul>
Latency	<ul style="list-style-type: none"> <li>■ Maximum network latency between NSX-T Managers is 10 ms.</li> <li>■ Maximum network latency between the NSX-T Manager cluster and transport nodes is 150 ms.</li> </ul>

## 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. The witness appliance should be running the same version of ESXi as the ESXi hosts in the stretched cluster. The maximum RTT on the witness is 200ms.

See [Deploy and Configure the vSAN Witness Host](#) in the VMware Validated Design document. This document follows a preset object naming convention. Treat this as an example and follow your own naming convention.

## Stretch a Cluster

This procedure describes how to stretch a cluster across two availability zones. Stretch clusters are supported only for VUM-based VI workload domains.

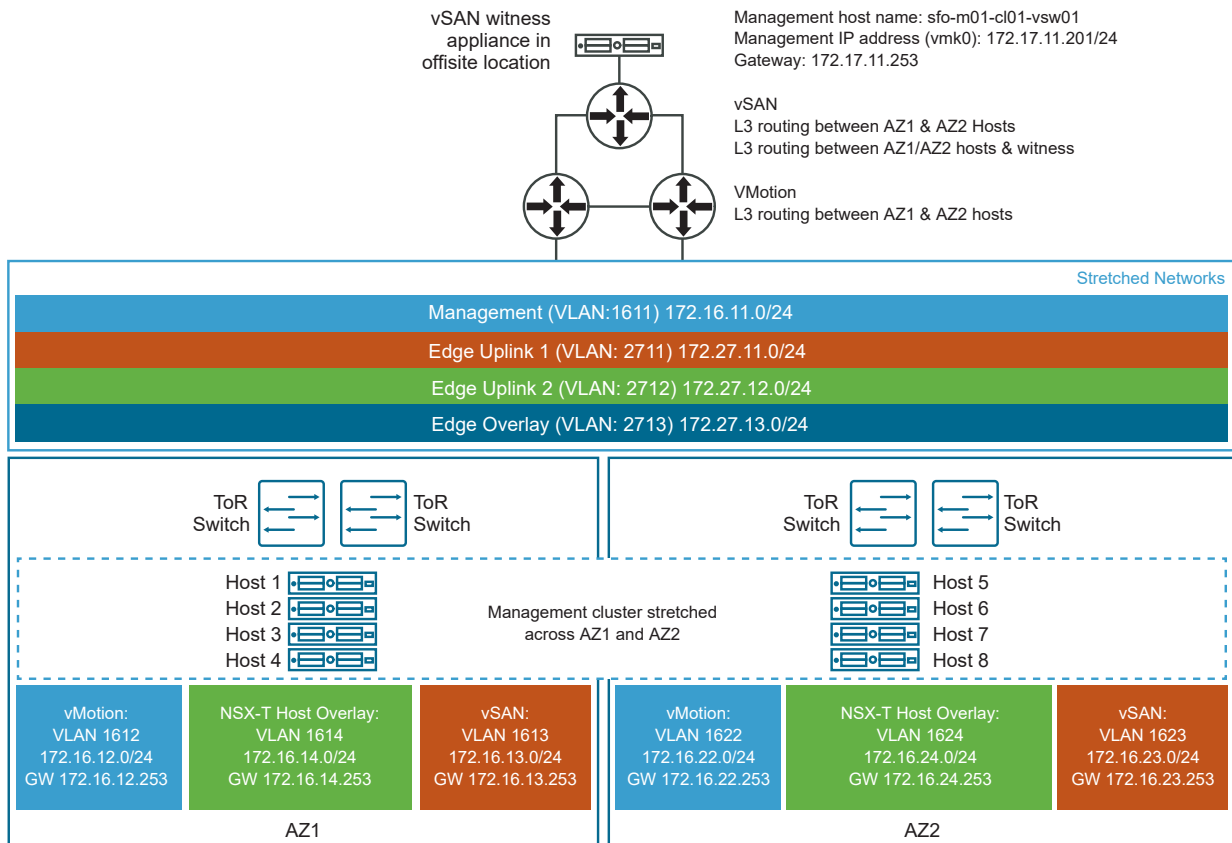
Our example use case has two availability zones in two buildings in an office campus - AZ1 and AZ2. Each availability zone has its own power supply and network. The management domain is on AZ1 and contains the default cluster, **SDDC-Cluster1**. This cluster contains four ESXi hosts. AZ1 also contains the default bring-up pool, bringup-networkpool.

vSAN network	VLANID=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	VLANID=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 Cloud Foundation inventory yet.

We will stretch the default cluster `SDDC-Cluster1` in the management domain from AZ1 to AZ2.

Figure 14-1. Stretch Cluster Example



## Prerequisites

- You must have deployed and configured a vSAN witness host. See [Deploy and Configure vSAN Witness Host](#).
- Configure routing to the datacenter infrastructure in the second availability zone and set the infrastructure VM restart order in case of a failure. See [NSX-T Data Center Configuration for Availability Zone 2](#) in the VMware Validated Design document. This document follows a preset object naming convention. Treat this as an example and follow your own naming convention.
- All VMs on an external network must be on an overlay backed segment. If they are on a VLAN, that VLAN must be stretched as well.
- 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.
- Ensure that the required TCP and UDP ports for vSAN network are open for communication between the availability zones and between the witness host and both availability zones. See KB article [52959](#).
- If you are stretching a cluster in a VI workload domain, the default management domain cluster must have been stretched.

## Procedure

- 1 Create a network pool, `AZ2-networkpool1`, on AZ2. See [Create a Network Pool](#).

Based on our example, here are the network details for the network pool.

vSAN network	VLANID=1623
	MTU=9000
	Network=172.16.234.0
	Netmask=255.255.255.0
	Gateway=172.16.234.1
	IP range= 172.16.234.11 - 172.16.234.59
vMotion network	VLANID=1612
	MTU=9000
	Network=172.16.22.0
	netmask 255.255.255.0
	gateway 172.16.22.1
	IP range= 172.16.22.11 - 172.16.22.59

- 2 Commission the four hosts in AZ2 and associate them with `AZ2-networkpool`. In our example, these are 172.16.21.105, 172.16.21.106, 172.16.21.107, 172.16.21.108.

See [Commission Hosts](#).

- 3 Get the UUIDs of the hosts in AZ2.
  - a On the SDDC Manager Dashboard, 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 UUIDs of the hosts in AZ2.
- 4 Get the ID of the cluster you are stretching (`SDDC-Cluster1` in our example).
  - 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 cluster ID for `SDDC-Cluster1`.



## 5 Validate the stretch cluster API.

### a Prepare the JSON request body.

- 1 In the API section for **Managing Clusters**, click `POST /v1/clusters/id/validations`.
- 2 Under `clusterUpdateSpec`, click `Cluster Update Data`  
`ClusterOperationSpecValidation`.
- 3 Click **Download** to download the JSON file.
- 4 Edit the downloaded JSON file so that it contains only the stretch section similar to the example below.

```
{
  "clusterUpdateSpec" : {
    "clusterStretchSpec": {
      "hostSpecs": [ {
        "id": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx",
        "licenseKey": "XXXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
      }, {
        "id": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx",
        "licenseKey": "XXXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
      }, {
        "id": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx",
        "licenseKey": "XXXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
      }, {
        "id": "xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxxxx",
        "licenseKey": "XXXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
      } ],
      "secondaryAzOverlayVlanId": 1624,
      "witnessSpec": {
        "fqdn": "sfo03m01vsanw01.sfo.rainpole.local",
        "vsanCidr": "172.17.13.0/24",
        "vsanIp": "172.17.13.201"
      }
    }
  }
}
```

- 5 Update the values for witness FQDN, CIDR, vSAN IP address, and host IDs and license keys.
- ### b Run the API command.
- 1 Click `POST /v1/clusters/id/validations`.
  - 2 For the `ClusterOperationSpecValidation` field, update the cluster UID (you retrieved this in step 4) and host UID ((you retrieved this in step 3).
- ### c Click **Execute**.

Ensure that the validation is successful. If there are any errors, resolve them and run the validation again.

## 6 Stretch the cluster.

### a Prepare the JSON request body.

- 1 Click `Patch /v1/clusters/id`.
- 2 Under `ClusterUpdateSpec`, click `Cluster Update Data`  
`ClusterUpdateSpec{ ... }`.
- 3 Click **Download** to download the JSON file.
- 4 Edit the JSON file so that it contains only the stretch section similar to the example below.

```
{
  "clusterStretchSpec": {
    "hostSpecs": [ {
      "id": "2c1744dc-6cb1-4225-9195-5cbd2b893be6",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
    }, {
      "id": "6b38c2ea-0429-4c04-8d2d-40a1e3559714",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
    }, {
      "id": "5b704db6-27f2-4c87-839d-95f6f84e2fd0",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
    }, {
      "id": "5333f34f-f41a-44e4-ac5d-8568485ab241",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX"
    } ],
    "secondaryAzOverlayVlanId": 1624,
    "witnessSpec": {
      "fqdn": "sfo03m01vsanw01.sfo.rainpole.local",
      "vsanCidr": "172.17.13.0/24",
      "vsanIp": "172.17.13.201"
    }
  }
}
```

- 5 Update the values for witness FQDN, CIDR, vSAN IP address, and host IDs and license keys.
- 6 Click **Execute**.  
The stretch cluster task is displayed in the SDDC Manager task panel.
- 7 Monitor the stretch cluster task till it is completed.

- 7 Log in to the vCenter Server to validate that the cluster has been stretched correctly.
  - a Check the vSAN health page.
    - 1 Click **Host and Clusters**.
    - 2 Select the stretched cluster (SDDC-Cluster1 in our example).
    - 3 Click **Monitor > vSAN > Health**.
    - 4 Click **Retest**.
    - 5 Fix errors, if any.
  - b Check the vSAN storage policy page.
    - 1 Click **Policies and Profiles > VM Storage Policies > vSANDefault Storage Policies**.
    - 2 Select the policy associated with the vCenter Server for the stretched cluster (SDDC-Cluster1 in our example)
    - 3 Click **Monitor > VMs and VirtualDisks**.
    - 4 Click **Refresh**.
    - 5 Click **Trigger VM storage policy compliance check**.
    - 6 Fix errors, if any.

#### What to do next

Add new VMs to the primary AZ and associate the host rule for the new VMs with the primary AZ rule.

## Expand a Stretched Cluster

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.

#### Procedure

- 1 Commission the additional hosts to Cloud Foundation.
 

See [Commission Hosts](#).
- 2 Get the UUIDs of the hosts you commissioned.
  - a On the SDDC Manager Dashboard, 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 UUIDs of the hosts.

- 3 Get the ID of the cluster you are expanding (SDDC-Cluster1 in our example).
  - 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 for SDDC-Cluster1.
- 4 Get the primary and secondary availability zone names from vCenter Server.
  - a Log in to vCenter Server.
  - b Navigate to **Cluster > Configure > vSAN > Fault Domains**.
  - c Note down the primary and secondary availability zone names.

## 5 Run the expand API command.

### a Prepare the JSON request body.

- 1 Click `Patch /v1/clusters/id`.
- 2 Under `ClusterUpdateSpec` field, click `Cluster Update Data`  
`ClusterUpdateSpec{ ... }`.
- 3 Click **Download** to download the JSON file.
- 4 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.

```
{
  "clusterExpansionSpec": {
    "hostSpecs": [ {
      "id": "2c1744dc-6cb1-4225-9195-5cbd2b893be6",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX",
      "azName": "primary/secondary"
    }, {
      "id": "6b38c2ea-0429-4c04-8d2d-40a1e3559714",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX",
      "azName": "primary/secondary"
    }, {
      "id": "5b704db6-27f2-4c87-839d-95f6f84e2fd0",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX",
      "azName": "primary/secondary"
    }, {
      "id": "5333f34f-f41a-44e4-ac5d-8568485ab241",
      "licenseKey": "XXXXX-XXXXX-XXXXX-XXXXX-XXXXX",
      "azName": "primary/secondary"
    } ]
  }
}
```

### b Run the expand cluster API.

- 1 For the `ClusterUpdateSpec` field, update the cluster ID (you retrieved this in step 3) and unstretch JSON file with the payload you prepared in step 5a.iii.
  - 2 Click **Execute**.  
The stretch cluster task is displayed in the SDDC Manager task panel.
  - 3 Monitor the task till it is completed.
- 6 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.

- 7 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 **Hosts and Clusters** and expand the stretched cluster (SDDC-Cluster1 in our example).
  - c Click the **Configure** tab.
  - d Under **Services**, click **vSphere Availability**, and then click **Edit**.
  - e On the Admission Control page of the Edit Cluster Settings dialog box, set **Host failures cluster tolerates** to the number of hosts in AZ1 and click **OK**.

## Unstretch a Cluster

This procedure describes how to unstretch a vSAN cluster.

In this procedure, we will unstretch the default cluster in the management domain (**SDDC-Cluster1**) that has been stretched from AZ1 to AZ2.

### Procedure

- 1 Get the ID of the cluster you are unstretching (SDDC-Cluster1 in our example).
  - a On the SDDC Manager Dashboard, 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 Unstretch the cluster.

### a Prepare the JSON request body.

- 1 Click `Patch /v1/clusters/id`.
- 2 Under `ClusterUpdateSpec` field, click `Cluster Update Data`  
`ClusterUpdateSpec{ ... }`.
- 3 Click **Download** to download the JSON file.
- 4 Edit the downloaded JSON file so that it contains only the unstretch information similar to the example below.

```
{
  "clusterUnstretchSpec": {}
}
```

### b Run the unstretch cluster API.

- 1 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 2a.
- 2 Click **Execute**.

The unstretch cluster task is displayed in the SDDC Manager task panel.

- 3 Monitor the unstretch cluster task till it is completed.

All hosts from AZ2 are removed from the unstretched cluster and 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.

### Prerequisites

- Image the replacement host with the same ESXi version as the other hosts in the cluster.
- Check the health of the cluster.

See "Check vSAN Health" in *Administering VMware vSAN*.

### Procedure

- 1 Get the ID of the host to be removed.
  - a On the SDDC Manager Dashboard, 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 Run the compact cluster API.
- a Prepare the JSON request body.
    - 1 Click `Patch /v1/clusters/id`.
    - 2 Under `ClusterUpdateSpec`, click **Cluster Update Data** `ClusterUpdateSpec{ ... }`.
    - 3 Click **Download** to download the JSON file.
    - 4 Edit the JSON file so that it contains only the compact section similar to the example below.

```
{
  " clusterCompactionSpec": {
    "hosts": [ {
      "id": "2c1744dc-6cb1-4225-9195-5cbd2b893be6"
    }, {
      "id": "6b38c2ea-0429-4c04-8d2d-40a1e3559714"
    }, {
      "id": "5b704db6-27f2-4c87-839d-95f6f84e2fd0"
    } ]
  }
}
```

- 5 In the `id` field, replace the values with the host IDs you retrieved in step 1.
  - 6 Click **Execute**.
- The compact cluster task is displayed in the SDDC Manager task panel.
- 7 Monitor the task till it is completed.
- 4 Decommission the host to be removed.
- See [Decommission Hosts](#).
- 5 Commission the replacement host to the same network pool as the removed host.
- See [Commission Hosts](#) .
- 6 Expand the cluster to add the commissioned host to the cluster. See [Expand a Stretched Cluster](#).



- 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.

#### **Results**

vSAN automatically rebuilds the stretch cluster.

# Composability Management

# 15

With composability, you can dynamically configure servers to meet the needs of your workloads without physically moving any hardware components. You bind disaggregated hardware components (compute, network, storage, and offload components) together to create a logical system based on the needs of your applications. These logical systems function like traditional rack mount systems.

The Cloud Foundation composability feature is available for HPE Synergy and Dell MX servers and uses the Redfish translation layer to connect to the composable hardware infrastructure. Redfish Translation Layer supports data models used to get composable resources and zones restrictions from the hardware infrastructure. It is designed to be extensible and vendor agnostic. You must obtain and install the Redfish appliance from the composable hardware vendor.

It is recommended that you compose and decompose servers only through Cloud Foundation, and not through the vendor software.

---

**Note** Ensure that you follow the restriction on storage sled placement on Dell MX servers. Refer to vendor documentation for more information.

---

This chapter includes the following topics:

- [Configure Translation Layer](#)
- [Compose a Server](#)
- [View Composability Information](#)
- [Add Storage](#)
- [Remove Storage](#)
- [Decompose a Server](#)

## Configure Translation Layer

Redfish translation layer is the interface between SDDC Manager and hardware vendor. You must configure this translation layer by providing the Redfish translation layer URL and credentials.

**Procedure**

- 1 As a best practice, increase the queue capacity for the thread pool.

- a Open the `application-prod.properties` file:

```
vi /opt/vmware/vcf/operationsmanager/config/application-prod.properties
```

- b Update the queue capacity line as follows:

```
om.executor.queuecapacity=300
```

- c Save and close the file.

- 2 If you are using a self-signed certificate, import the Redfish certificate from the Redfish VM to SDDC Manager VM by following the steps below. If you are using a CA signed certificate, skip to step 3.

- a Using SSH, log in to the SDDC Manager VM with the following credentials:

User name: `vcf`

Password: use the password specified in the deployment parameter workbook.

- b Enter `su` to switch to the `root` user.

- c Import the Redfish certificate from the Redfish VM to SDDC Manager VM by running the following command:

```
/opt/vmware/vcf/commonsvcs/scripts/cert-fetch-import-refresh.sh --ip=redfish-ip --
port=port --service-restart=operationsmanager

ip Specify translation layer IP
port TLS/SSL port
```

The output displays information about the certificate to import including owner, issuer, serial number, validity, certificate fingerprints (md5, sha1, or sha256), signature algorithm name, subject, public key algorithm, and version. Verify this information.

- d Answer the prompt.

Operations Manager is restarted. Wait for a few minutes and then proceed to step 4.

- 3 Restart Operations Manager:

```
systemctl restart operationsmanager
```

Wait for a few minutes before proceeding to the next step.

- 4 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.

- 5 Enter the URL for the Redfish translation layer.

- 6 Enter the user name and password for the Redfish translation layer.

- 7 Click **Connect**.

# Compose a Server

You can compose one or more servers by selecting the compute, network, and storage resources.

## Prerequisites

- The translation layer must have been configured.
- The composed server must meet the minimum hardware requirements. See the *Planning and Preparation Workbook*.

## Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 2 From the Available Resources table, select the zone where you want to compose a server. A zone corresponds to a physical boundary.
- 3 Click **Compose**.
- 4 In the Allocate Resources dialog box, select the compute for the server.
- 5 Select the storage.
- 6 Select the network interface.

The Choose number of servers section displays the number of servers you can compose based on the selected resources.

- 7 Select the number of servers you want to compose.
- 8 Click **Next**.
- 9 On the Review page, review the allocated resources to the servers.  
Click **Back** to make any changes.
- 10 Click **Compose**.

## Results

The compose server task is displayed in the Tasks table at the bottom of the Composable Infrastructure page. Click the name of the task for more information. When the server is composed, it is added to the Server Composition Summary table.

If Cloud Foundation does not receive a response from Redfish due to an external error, an error message is displayed. The hardware resources in the compose request are locked and cannot be used. You must free up these resources using the vendor UI. For more information, refer to the vendor documentation.

## What to do next

- 1 Image the composed servers. See [Chapter 6 Installing ESXi Software on Cloud Foundation Servers](#).
- 2 Commission the composed servers. See [Commission Hosts](#).

## View Composability Information

The Composable Infrastructure page displays information about available resources and composed servers.

### Procedure

- ◆ On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.

The Composable Infrastructure page appears. The Redfish translation layer information is displayed on the top of the page.

The Available Resources table displays the available zones and computer, storage, and network information available in each zone.

The Server Composition Summary table displays the composed servers.

The task panel at the bottom of the page shows the tasks performed and their status.

## Add Storage

You can add storage to Dell MX composed servers.

### Prerequisites

The server you are adding storage to must be unassigned.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 2 In the Composed Servers section, click the three dots next to the server you want to add storage to and click **Add Storage**.  
  
The Allocate Storage dialog box displays the composed server details. Hovering over the Information icon next to the Current Storage field displays the SSD and HDD details of the server.
- 3 In the Add to Current Storage section, select the storage units you want to add and click **Next**.
- 4 On the Review page, verify the storage you are adding and click **Add**.

## Remove Storage

You can remove storage units from a Dell MX composed server.

### Prerequisites

The server you are removing storage from must be unassigned.

### Procedure

- 1 On the Composable Infrastructure page, click the three dots next to the server you want to remove storage from.
- 2 Click **Remove Storage**.
- 3 For the storage type you want to remove, set the Remove from Allocated value to the appropriate value.
- 4 Click **Next**.
- 5 On the Review page, verify the storage you are removing and click **Remove**.

## Decompose a Server

You can decompose a server that has not been assigned to a VI workload domain.

### Prerequisites

The server to be decomposed must be unassigned.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Composable Infrastructure**.
- 2 From the Server Composition Summary table, select the server to be decomposed.
- 3 Click **Decompose**.
- 4 In the Decompose Servers dialog box, click **Decompose**.

# Monitoring Capabilities in the Cloud Foundation System

# 16

The 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.

The monitoring capabilities involve these features:

This chapter includes the following topics:

- [Viewing Tasks and Task Details](#)

## Viewing Tasks and Task Details

From the SDDC Manager Dashboard, 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 the SDDC Manager Dashboard, see [Tour of the SDDC Manager 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.



# Updating Cloud Foundation DNS and NTP Servers

# 17

If you need to make changes to the DNS or NTP server information that you provided during Cloud Foundation bring-up, you can use the VMware Cloud Foundation API to update the servers.

When you initially deploy 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 change this server information at a later date, using the VMware Cloud Foundation API.

This chapter includes the following topics:

- [Update DNS Server Configuration](#)
- [Update NTP Server Configuration](#)

## Update DNS Server Configuration

Use this procedure to update the DNS server information that you provided during Cloud Foundation bring-up.

Cloud Foundation uses DNS servers to provide name resolution for various components in the system. You must provide root DNS domain information. Optionally, you can provide subdomain information. When you update the DNS server configuration, Cloud Foundation updates the components in a specific order:

- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX-T Managers
- vRealize Suite Lifecycle Manager

---

**Note** This procedure does not update NSX-T Edge nodes.

---

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 vRealize Suite Lifecycle Manager. 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 Cloud Foundation API, which is secured by token-based authentication.

### Prerequisites

- Ensure that both forward and reverse DNS resolution is functional for each component using the updated DNS server.
- The new DNS server should be reachable from the Cloud Foundation components.
- All Cloud Foundation components should be reachable from SDDC Manager.
- All Cloud Foundation components must be in an `Active` state.

## Procedure

- 1 To obtain an access token, run the following command:

```
curl 'https://SDDC_MANAGER_IP/v1/tokens' -k -X POST -H 'Content-Type: application/json' -H 'Accept: application/json' -d '{"username": "user name", "password": "user password"}'
```

Replace the SDDC Manager IP address, user name, and password with the information for your environment.

The command returns an access token and a refresh token.

[illegible]

- 2 Get the current DNS server configuration information.

[illegible]

- 3 Validate the new DNS server configuration information.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUC00NGQwLTQ5MzYtYjQwMC0xMzc5NzMyMjdmOWUiLCJpYXQiOjE1ODU3ODQ5MTg5InN1YiI6ImFkbWl1aXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsImVzcyI6InZjZilhdXRoIiwiaXYXVkiJoic2RkYyIzZXJ2aWNlcysIm5izi
```

```
I6MTU4NTc4NDkxOCwiZXhwIjoxNTg1Nzg4NTE4LCJlc2VyIjoiYWRtaW5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbCIsInNjb3B1IjpbIkJBQ0tVUF9DT05GSUdfUkVBRClSikNSRURFTlRJRQ0xfUkVBRClSIlVTRVJfV1JJVEUilCJpVEhFUl9XUk1URSlSikJBQ0tVUF9DT05GSUdfV1JJVEUilCJpVEhFUl9SRUFEIiwiVWVFNl9SRUFEIiwiQ1JFREVOVElBTF9XUk1URSlSIdfQ. Ya4XsZntsRHUZFRBNKGy7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X POST https://SDDC_MANAGER_IP/v1/system/dns-configuration/validations -d '{"dnsServers":[{"ipAddress":"<dns-server-ip>","isPrimary":true}]}' | json_pp
```

Replace `<dns-server-ip>` with the IP address of the new DNS server. Specify `true` or `false` for `isPrimary`, depending on whether or not the new DNS server is the primary DNS server.

Note the *<id>* that gets returned.

The validator verifies forward and reverse name resolution for Cloud Foundation components using the new DNS server.

- 4 Monitor the status of the validation task.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxNmM5YThmOC00NGQwLTQ5MzYtYjJqWmC0xMzc5NzMzMjdmOWUiLCJpYXQiOiE1ODU3ODQ5MTgsInNlYyI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsImVycyI6InZjZilhdXRoIiwiaXYXkKIjoic2RkYy1zZXJ2aWNlcysIm5iziI6MTU4NTc4NDkxOClwZWxzXhwIjoxtNgTlNzg4NTE4LGljc1c2VyIjoyYWRTaw5pc3RyYXRvckB2c3BoZXJlLmxxy2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsInNjb3BlIjpjbIkKBQ0tVUF9DT05GSudfUkVBRCIsIkNSRURFTlRJQUxfukVBRCIsIlVTRVFvfv1JJVEUilCJPVEhfUL9XUKlURSIsIkJBQ0tVUF9DT05GSudfv1JJVEUilCJPVEhfUL9SRUFEIiwiVVNFUL9SRUFEIiwiQ1JFREVOVELBTFF9XUKlURSIdfQ.Ya4XSzntsRHUZFRBNKGY7Js6xrGYGe8KgDj2QBihFmg" -H 'Content-Type: application/json' -k -X GET https://SDCC_MANAGER_IP/v1/system/dns-configuration/validations/<id>ljson_pp
```

Replace `<id>` with the ID from the previous step.

If validation succeeds, you can proceed to change the DNS server configuration. If validation fails, correct any issues and try again.

**5** Change the DNS server configuration information.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxNzM5YThmOC00NGQwLTQ5MzYtYTljZWMC0xMzc5NjMyMjdmdWUiLCJpYXQiOjE1ODU3ODQ5MTgsInNlYyI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCIslmlzcyI6InZjZilhdXRoIiwiaXYXVkaioic2RkYylzZXZjaWNlcysIm5iziI6MTU4NTc4NDkxOWCwiZhXhwIjoxtNgTglnZzg4NTE4LCljc2VyIjoyYWRTaW5pc3RyYXRvcKB2c3BoZXJlLmxxy2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5lb2NhbmCIslbnjb3BlIjpjbIkJBQ0tvUF9DT05GSudfUkVBRCIsIkNSRUFRFTlrJQUxfukVBRCIsIlVTRVJfv1JJVEUiLCJPVEhfUL9XUKlURSlSkJBQ0tvUF9DT05GSudfv1JJVEUiLCJPVEhfUL9SRUFEIiwiVVNFUL9SRUFEIiwlQ1JFREVOVELBTFF9XUKlURSIdjfQ.Ya4xsZntsRHUZFRBNKGY7Js6xrGYGe8KgDJ2QBihFmg" -H 'Content-Type: application/json' -k -X PUT https://SDDC_MANAGER_IP/v1/system/dns-configuration -d '{"dnsServers":[{"ipAddress":"<dns-server-ip>", "isPrimary":true}]}' |json pp
```

Replace `<dns-server-ip>` with the IP address of the new DNS server. Specify `true` or `false` for `isPrimary`, depending on whether or not the new DNS server is the primary DNS server.

Note the *<id>* that gets returned.

## 6 Track the status of the DNS update.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxN2M5YThmOC00NGQwLTQ5MzYtYjYwMC0xMzc5N2MyMjdmOWUwLjIjYXQ5MTgsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsImVzcyI6ImZlhdXRoIiwiaXVzYXki
```

```
joic2RkYy1zZXJ2aWNlcysIm5iZiI6MTU4NTc4NDkxOCwiZXhwIjoxNTg1Nzg4NTE4LCJ1c2VyIjoiYWRtaW5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsInNjb3BlIjpbIkJBQ0tVUF9DT05GSUdfUkVBRCIsIkNSRURFTlRjQUxfUkVBRCIsIlVTRVJfV1JJVEUiLCJpVEhFU19XUklURSIsIkJBQ0tVUF9DT05GSUdfV1JJVEUiLCJpVEhFU19SRUFEIiwVNVF19SRUFEIiwQ1JFREVOVE1BTF9XUklURSJdfQ.Ya4XsZntsRHUZFBNKGY7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X GET https://SDDC_MANAGER_IP/v1/tasks/<id>|json_pp
```

Replace *<id>* with the ID from the previous step. Wait for the task to complete.

## 7 Verify that the DNS configuration was updated.

```
curl 'https://SDDC_MANAGER_IP/v1/system/dns-configuration' -k -X GET -H "Content-Type: application/json" -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiI2ZGRmZmU3MC0zZGEwLTQ5NmUtOTQzMCIhOGNkNTQ0YTk2ZGMiLCJpYXQiOiJlODU3ODE3ODcsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsImZlcyI6InZjZilhdXRoiIiwiaXVkiOiIjoic2RkYy1zZXJ2aWNlcysIm5iZiI6MTU4NTc4MTc4NywiZXhwIjoxNTg1Nzg1Mzg3LCJ1c2VyIjoiYWRtaW5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsInNjb3BlIjpbIkJBQ0tVUF9DT05GSUdfUkVBRCIsIkNSRURFTlRjQUxfUkVBRCIsIlVTRVJfV1JJVEUiLCJpVEhFU19XUklURSIsIkJBQ0tVUF9DT05GSUdfV1JJVEUiLCJpVEhFU19SRUFEIiwVNVF19SRUFEIiwQ1JFREVOVE1BTF9XUklURSJdfQ.WCpUPRIm5A6X_406HTJF7TbTSa0g91_AQbt70cBPblM"
```

## Update NTP Server Configuration

Use this procedure to update the NTP server information that you provided during Cloud Foundation bring-up.

Cloud Foundation uses NTP servers to synchronize time between the various components in the system. You must have at least one NTP server. When you update the NTP server configuration, Cloud Foundation updates the components in a specific order:

- SDDC Manager
- vCenter Servers
- ESXi hosts
- NSX-T Managers
- vRealize Suite Lifecycle Manager

---

**Note** This procedure does not update NSX-T Edge nodes.

---

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 vRealize Suite Lifecycle Manager. 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 Cloud Foundation API, which is secured by token-based authentication.

### Prerequisites

- Any new NTP server is reachable by all components.
- Time skew between new NTP servers is less than 5 minutes.

## Procedure

- 1 To obtain an access token, run the following command:

```
curl 'https://SDDC_MANAGER_IP/v1/tokens' -k -X POST -H 'Content-Type: application/json' -H 'Accept: application/json' -d '{"username" : "user_name", "password" : "user_password"}'
```

Replace the SDDC Manager IP address, user name, and password with the information for your environment.

The command returns an access token and a refresh token.

[illegible]

- 2 Get the current NTP server configuration information.

[illegible]

- 3 Validate the new NTP server configuration information.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUc00NGQwLTQ5MzYtYjQwMC0zMzc5NzMyMjdmOWUiLCJpYXQiOjE1ODU3ODQ5MTg3InR1YiI6ImFkbWlwuaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCIslmVycyI6InZjZilhdXRoIiwiaXYxvkiJoic2RkYyIzZXJ2aWNlcyIsIm5iZiI6MTU4NTc4NDKxOCwiZXBhIjoxtNg4NTE4LClc2VyIjoyYWRTaW5wc3RyYXRvcB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCIslmNjb3BlIjpjbikJBQ0tVUF9DT05GSUdfUkVBRClSlkNSRURFTlRJQUxfUkVBRClSlVTRVJfVlJJVEUilCJpVEhfUl9XuklURSlSkJBQ0tVUF9DT05GSUdfVlJJVEUilCJpVEhfUl9SRUFEIiwiVVNFUl9SRUFEIiwiQ1JFREVOVELBTBF9XuklURSJdfQ.Ya4XszntsRHUZFRBNKGy7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X POST https://SDDC_MANAGER_IP/v1/system/ntp-configuration/validations -d '{"ntpServers":[{"ipAddress":"<ntp_server-ip>}"]}' | json_pp
```

Replace `<ntp-server-ip>` with the IP address of the new NTP server. To enter multiple NTP servers, use the following format: `{"ntpServers":[{"ipAddress":"<ntp-server-ip-1>"}, {"ipAddress":"<ntp-server-ip-2>"}]}`.

Note the *<id>* that gets returned.

The validator verifies that the Cloud Foundation components can communicate with the new NTP server.

4 Monitor the status of the validation task.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxNm5YThmOC00NGQwLTQ5MzYtYjJqWmC0xMzc5NzMyMjdmOWUiLCJpYXQiOiE1ODUzODQ5MTgsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbcIsImVycyI6InZjZilhdXRoIiwiaXVkaWoiOiRkYylyZXZlZWNLcyIsIm5iziI6MTU4NTc4NDkxOCwiZXhwIjoxtNg4NTE4LCJlc2VyIjoieWRtaW5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5lb2NhbcIsInNjb3BlIjpjbIkjBQ0tVUF9DT05GSUdfUkVBRCIsIkNSRURFTlRJQUxfUkVBRCIsIlVTRVJfvlJJVEUiLCJPVEhfU19XUKlURSIsIkJBQ0tVUF9DT05GSUdfv1JJVEUiLCJPVEhfU19SRUFEIiwivVNvF19SRUFEIiw1JFREVOVElBTf9XUklURSIdfQ.Ya4XsZntsRHUZFRBNKGY7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X GET https://SDCC_MANAGER_IP/v1/system/ntp-configuration/validations/<id>|json_pp
```

Replace `<id>` with the ID from the previous step.

If validation succeeds, you can proceed to change the NTP server configuration. If validation fails, correct any issues and try again.

**5** Change the NTP server configuration information.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxNzMyYThmOC00NGQwLTQ5MzYtYjJwMC0xMzc5NzMyMjdmOWUiLCJpYXQiOjE1ODUzODQ5MTgsInNlYiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCIslmZcyI6InZjZilhdXRoIiwiaXVkaWoiOiRkYy1zZXZlZWNLcyIsIm5iZiI6MTU4NTc4NDkxOCwiZXhwIjoxtNgTlNzg4NTE4LClc2VyIjoiyWRtaW5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCIslmNjb3BlIjpjbIkjBQ0tVUF9DT05GSudfUkVBRCIsIkNSRURFTlRJRQ0xUkVBRCSIlVTRVJfv1JJVEUilCJPVEhfUL9XUklURSI6IkjBQ0tVUF9DT05GSudfV1JJVEUilCJPVEhfUL9SRUFEIiwivVNfUL9SRUFEIiwq1JFREVOVELBTf9XUklURSIdfQ.Ya4XsZntsRHUZFRBNKGY7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X PUT https://SDDC_MANAGER_IP/v1/system/ntp-configuration -d '{"ntpServers":[{"ipAddress":"<ntp-server-ip>}]}' | json_pp
```

Replace `<ntp-server-ip>` with the IP address of the new NTP server. To enter multiple NTP servers, use the following format: `{"ntpServers":[{"ipAddress":"<ntp-server-ip-1>"}, {"ipAddress":"<ntp-server-ip-2>"}]}`.

Note the *<id>* that gets returned.

## 6 Track the status of the NTP update.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxN2M5YThmOC00NGQwLTQ5MzYtYjQwMC0xMzc5N2MyMjdmOWUwLjIjYXQlOiJlODU3ODQ5MTgsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCI9ImVzcyI6InZjZi1hdXR0eS9iYXVxIjoic2RkYy1zZXZlZW50IiwiaWF0IjoiMTU0NTE4NDkxOCwiZG90IjojNzZlNTE4LCJlc2VyIjojYWRtaW50c3RyYXRvcB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbmCI9InNjb3BlIjpbIkU
```

```
BQ0tVUF9DT05GSUdfUkVBRCIsIkNSRURFTlRjQUxfUkVBRCIsIlVTRVJfVlJJVEUiLCJPVEhFUl9XUklURSIkJBQ0
tVUF9DT05GSUdfVlJJVEUiLCJPVEhFUl9SRUFEIiwiVVNFUl9SRUFEIiwiQ1JFREVOVElBTF9XUklURSJdfQ.Ya4XsZ
ntsRHUZFBNKGY7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X GET
https://SDDC_MANAGER_IP/v1/tasks/<id>|json_pp
```

Replace *<id>* with the ID from the previous step. Wait for the task to complete.

## 7 Verify that the NTP configuration was updated.

```
curl 'https://SDDC_MANAGER_IP/v1/system/ntp-configuration' -k -X GET -H "Content-Type:
application/json" -H "Authorization: Bearer
eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiI2ZGRmZmU3MC0zZGEwLTQ5NmUtOTQzMCIhOGNkNTQ0YTk2ZGMiLCJpYXQiOiJ
E1ODU3ODE3ODcsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbCIsImIzcyI6InZjZi1hdXRoIiwiaXVki
joic2RkYy1zZXJ2aWNlcysIm5iZiI6MTU4NTc4MTc4NywiZXhwIjoxNTg1NzglMzg3LCJlc2VyIjoieYWRtaW5pc3Ry
YXRvc2B2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbCIsInNjb3BlIjpbIk
BQ0tVUF9DT05GSUdfUkVBRCIsIkNSRURFTlRjQUxfUkVBRCIsIlVTRVJfVlJJVEUiLCJPVEhFUl9XUklURSIkJBQ0
tVUF9DT05GSUdfVlJJVEUiLCJPVEhFUl9SRUFEIiwiVVNFUl9SRUFEIiwiQ1JFREVOVElBTF9XUklURSJdfQ.WCpUPR
Im5A6X_406HTJF7TbTSa0g91_AQbt7OcBPblM"
```

# Supportability and Serviceability (SoS) Utility

# 18

The SoS utility is a command-line tool that you can use to run health checks, collect logs for Cloud Foundation components, and so on.

To run the SoS utility, SSH in to the SDDC Manager VM 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.

This chapter includes the following topics:

- [SoS Utility Options](#)
- [Collect Logs for Your Cloud Foundation System](#)

## SoS Utility Options

This section lists the specific options you can use with the SoS utility.

### 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 VMware Cloud Foundation Summary Options

These options provide information about the Cloud Foundation system 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
--get-vcf-summary	Returns information about your Cloud Foundation system, including CEIP, domains and clusters, hosts, licensing, network pools, SDDC Manager, and VCF services.
--get-vcf-services-summary	Returns information about SDDC Manager uptime and when Cloud Foundation services (for example, LCM) started and stopped.
--get-vcf-tasks-summary	Returns information about Cloud Foundation tasks, including the time the task was created and the status of the task.

## 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.

**Note** For generic options related to log collection, see [Collect Logs for Your Cloud Foundation System](#).

Option	Description
--configure-sftp	Configures SFTP for logs.
--debug-mode	Runs the SoS utility in debug mode.

Option	Description
<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> .  <b>Note</b> If you omit the <code>--domain-name</code> flag and domain name, the SoS operation is performed only on the management domain.
<code>--force</code>	Allows SoS operations to be formed while workflows are running.  <b>Note</b> It is recommended that you do not use this option.
<code>--history</code>	Displays the last 20 SoS operations performed.
<code>--ondemand-service</code>	Include this flag to execute commands on all ESXi hosts in a domain.  <b>Warning</b> Contact VMware 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 domain. For example, <code>/opt/vmware/sddc-support/&lt;JSON file name&gt;</code>
<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 VM in the <code>/opt/vmware/sddc-support/</code> directory.
<code>--skip-known-host-check</code>	Skips the specified check for SSL thumbprint for host in the known host.
<code>--zip</code>	Creates a zipped TAR file for the output.

## SoS Utility Options for Health Check

These SoS commands are used for checking the health status of various components or services, including connectivity, compute, storage, database, 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>--json-output-dir JSONDIR</code>	Outputs the results of any health check as a JSON file to the specified directory, <code>JSONDIR</code> .
<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 Managers, and SDDC Manager can be pinged.

Option	Description
<code>--composability-infra-health</code>	Performs an API connectivity health check of the composable infrastructure. If no composable infrastructure exists, this flag is ignored. If found, the utility checks connectivity status through the composable infrastructure API, such as Redfish.
<code>--compute-health</code>	Performs a compute health check, including ESXi host licenses, disk storage, disk partitions, and health status.
<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>--get-host-ips</code>	Returns host names and IP addresses of ESXi hosts.
<code>--get-inventory-info</code>	Returns inventory details for the Cloud Foundation components, such as vCenter Server NSX, SDDC Manager, and ESXi hosts. Optionally, add the flag <code>--domain-name ALL</code> to return all details.
<code>--hardware-compatibility-report</code>	Validates ESXi hosts and vSAN devices and exports the compatibility report.
<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 SDDC Manager VM. It also ensures that the hardware and software time stamp of ESXi hosts are within 5 minutes of the SDDC Manager VM.
<code>--password-health</code>	Returns the status of all current passwords, such as Last Changed Date, Expiry Date, and so on.
<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>--storage-health</code>	Performs a check on the vSAN disk health of the ESXi hosts and vCenter clusters. Also runs proactive vSAN tests to verify the ability to create VMs within the vSAN disks.
<code>--run-vsan-checks</code>	Runs proactive vSAN tests to verify the ability to create VMs within the vSAN disks.

## SoS Utility Options for Managing ESXi Hosts

Use these options to manage ESXi hosts, 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
```

Option	Description
<code>--disable-lockdown-esxi</code>	<p>Deactivate lockdown mode on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> <li>■ To deactivate lockdown on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>.</li> <li>■ To deactivate lockdown on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>.</li> </ul> <p><b>Note</b> If you do not specify domain, this command affects only the MGMT domain by default.</p>
<code>--enable-lockdown-esxi</code>	<p>Enables lockdown mode on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> <li>■ To enable lockdown on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>.</li> <li>■ To enable lockdown on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>.</li> </ul> <p><b>Note</b> If you do not specify domain, this command affects only the MGMT domain by default.</p>
<code>--disable-ssh-esxi</code>	<p>Deactivate SSH on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> <li>■ To deactivate SSH on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>.</li> <li>■ To deactivate SSH on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>.</li> </ul> <p><b>Note</b> If you do not specify domain, this command affects only the MGMT domain by default.</p>
<code>--enable-ssh-esxi</code>	<p>Enables SSH on ESXi nodes in the specified domains.</p> <ul style="list-style-type: none"> <li>■ To enable SSH on ESXi nodes in a specific domain, include the flag <code>--domain-name DOMAINNAME</code>.</li> <li>■ To enable SSH on ESXi nodes in all domains, include the flag <code>--domain-name ALL</code>.</li> </ul> <p><b>Note</b> If you do not specify domain, this command affects only the MGMT domain by default.</p>
<code>--refresh-ssh-keys</code>	Refreshes the SSH keys.

## SoS Utility Options for vRealize Suite Lifecycle Manager

Use these options to redeploy vRealize Suite Lifecycle Manager and monitor the redeployment. 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** You should only redeploy vRealize Suite Lifecycle Manager when directed to do so by VMware Support.

Option	Description
<code>--vrs lcm-redeploy</code>	Redeploys vRealize Suite Lifecycle Manager. Provides a taskID for the operation.
<code>--get-vrs lcm-redeploy-task-status &lt;taskID&gt;</code>	Returns vRealize Suite Lifecycle Manager redeployment status for the specified taskID.

## Collect Logs for Your 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.

- To collect all logs from all components, you can run the SoS utility without specifying any component-specific options.

- 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 Cloud Foundation environment.

**Table 18-1. SoS Utility Log File Options**

Option	Description
<code>--api-logs</code>	Collects output from REST endpoints for SDDC Manager inventory and LCM.
<code>--esx-logs</code>	Collects logs from the ESXi hosts only. Logs are collected from each ESXi host available in the deployment.
<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>--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>--no-health-check</code>	Skips the health check executed as part of log collection.
<code>--nsx-logs</code>	Collects logs from the NSX Manager and NSX Edge instances only.

Table 18-1. SoS Utility Log File Options (continued)

Option	Description
<code>--rvc-logs</code>	<p>Collects logs from the Ruby vSphere Console (RVC) only. RVC is an interface for ESXi and vCenter.</p> <p><b>Note</b> If the Bash shell is not enabled in vCenter, RVC log collection will be skipped .</p> <p><b>Note</b> RVC logs are not collected by default with <code>./sos</code> log collection. You must enable RVC to collect RVC logs.</p>
<code>--sddc-manager-logs</code>	Collects logs from the SDDC Manager only. <code>sddc&lt;timestamp&gt;.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>--test</code>	Collects test logs by verifying the files.
<code>--vc-logs</code>	<p>Collects logs from the vCenter Server instances only.</p> <p>Logs are collected from each vCenter server available in the deployment.</p>
<code>--vm-screenshots</code>	Collects all VM screenshots.
<code>--vrealize-logs</code>	Collects logs from vRealize Suite Lifecycle Manager.

## Procedure

- Using SSH, log in to the SDDC Manager VM with the following credentials:  
 Username: **vcf**  
 Password: use the password specified in the deployment parameter workbook.
- 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 VM

## Results

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.

## Example

```
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]
```

## What to do next

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>	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 <a href="https://nsxmanagerIPAddr/api/1.0/appliance-management/techsupportlogs/NSX">https://nsxmanagerIPAddr/api/1.0/appliance-management/techsupportlogs/NSX</a> , 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	Standard NSX Edge support bundle, generated using the NSX API to query the NSX Edge support logs: GET <a href="https://nsxmanagerIPAddr/api/4.0/edges/edgeId/techsupportlogs">https://nsxmanagerIPAddr/api/4.0/edges/edgeId/techsupportlogs</a> , 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.
<b>Note</b> This information is only collected if NSX Edges are deployed.	

## 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

# 19

This section provides procedures for repairing or replacing hosts in 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 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 VMware 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.

This chapter includes the following topics:

- [Avoiding Unintentional Downtime](#)
- [Replacing Components of a Host Running in Degraded Mode](#)
- [Replace a Dead Host](#)
- [Replace Boot Disk on a Host](#)

## 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 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 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 degraded depend on whether the host is assigned or unassigned. An assigned host belongs to the management domain or a workload domain. An unassigned host has been commissioned, but is not assigned to a domain.

These procedures apply to the following components:

- CPU
- Memory
- BMC
- Power supply

### Prerequisites

Before proceeding, review the guidance provided in the [Avoiding Unintentional Downtime](#).

## 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.

### Prerequisites

- Verify that the host is operational and is accessible by 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.

### Procedure

- 1 Log in to the vSphere Client.
- 2 Right-click the affected host and click **Maintenance Mode > Enter Maintenance Mode**.
- 3 If the host belongs to a domain, click **Full Data Migration**.
- 4 Right-click the affected 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 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.

### Prerequisites

- 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.

### Procedure

- 1 Log in to vSphere Web Client.
- 2 Right-click the affected host and select **Shutdown**.
- 3 Pull the host out of the physical rack.  
Note the ports on the switches it was connected to.
- 4 Service the appropriate part following the OEM vendor documentation.
- 5 Put the host back in the physical rack and connect it back to the appropriate switches.
- 6 Power on the host.
- 7 In the SDDC Manager Dashboard, 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.

This procedure applies chiefly to the following components:

- Storage controllers
- Motherboards
- Boot disks

### Prerequisites

If the host is assigned to a workload domain, verify that there are at least four hosts in the management or workload domain to which the faulty host belongs. If there are fewer than 4 hosts, contact VMware Support for assistance. Before proceeding, review the guidance provided in [Avoiding Unintentional Downtime](#).

### Procedure

- 1 If the host is assigned to a workload domain, it must be forcibly removed.
- 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.

### Prerequisites

If the host is operational, verify that there are at least four hosts in the management or workload domain to which the faulty host belongs. If there are fewer than four hosts, add a host to the domain from the capacity pool, if possible. If the host is not operational, see [Replace a Dead Host](#).

### Procedure

- 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](#).

# User and Group Management

# 20

You can allow the users and groups in your Microsoft Active Directory (AD) domain to use their credentials to log in to the SDDC Manager Dashboard as well as the vCenter Server instances that are deployed in your Cloud Foundation system.

You provided a password for the superuser account (user name `vcf`) in the deployment parameter workbook before bring-up. After Cloud Foundation is deployed, you can log in with the superuser credentials and then add vCenter Server or AD users or groups to Cloud Foundation. Authentication to the SDDC Manager Dashboard uses the VMware vCenter® Single Sign-On authentication service that is installed during the bring-up process for your Cloud Foundation system.

Users and groups can be assigned roles to determine what tasks they can perform from the UI and API.

This chapter includes the following topics:

- [Add a User or Group to Cloud Foundation](#)
- [Remove a User or Group](#)
- [Create a Service Account and Generate an Access Token](#)

## Add a User or Group to Cloud Foundation

You can add users or groups so that they can log in to SDDC Manager with their AD credentials.

### Prerequisites

Only a user with the ADMIN role can perform this task.

### Procedure

- 1 Log in to the SDDC Manager Dashboard with your superuser credentials.
- 2 Click **Administration > Users**.
- 3 Click **+ User or Group**.
- 4 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.

- 5 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.

- 6 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 Dashboard.

### Prerequisites

Only a user with the ADMIN role can perform this task.

### Procedure

- 1 On the SDDC Manager Dashboard, click **Administration > Users**.
- 2 Hover your mouse over the row containing the user or group that you want to remove.  
Three dots appear to the left of the user/group name column.
- 3 Click the dots and click **Remove**.

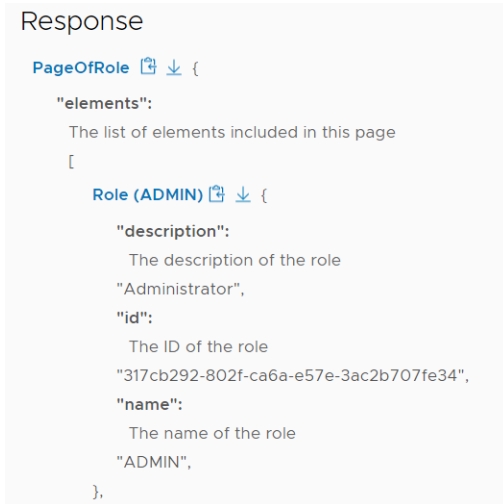
## Create a Service Account and Generate an Access Token

Cloud Foundation APIs are secured using token-based authentication. Create a service account and generate an access token so you can programmatically access Cloud Foundation APIs from other applications, and access the Cloud Foundation API when the management vCenter Server is down.

### Procedure

- 1 Log in to the SDDC Manager Dashboard as a user with the ADMIN role.  
For more about roles, see [Chapter 20 User and Group Management](#).
- 2 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**.

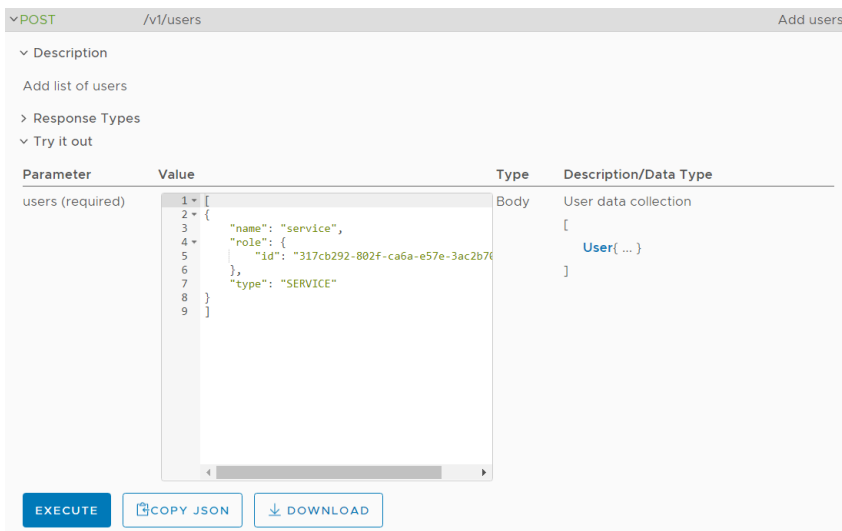
- c In the Response, click `PageOfRole` and `Role (ADMIN)`.
- d Copy the ID for the ADMIN role.



- 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:



Paste the ADMIN role ID from step 3.



- 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
          "qsfqnYgyxXQ892Jk90HXyuEMgE3SgfTS",

```

- 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.

The screenshot shows the API console for the endpoint `POST /v1/tokens` (Create Token Pair). The description states: "Creates access token and refresh token for user access". Under "Try it out", a table shows the configuration for the `tokenCreationSpec (required)` parameter:

Parameter	Value	Type	Description/Data Type
tokenCreationSpec (required)	<pre>1 = { 2   "apiKey": "qsfqnYgyxXQ892Jk90HXyuEMgE3SgfTS" 3 }</pre>	Body	tokenCreationSpec <a href="#">TokenCreationSpec{ ... }</a>

At the bottom, there are buttons for **EXECUTE**, **COPY JSON**, and **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  
    "eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiJkMWQxOGM5NiItZmZhLTQwODMtYjUzYiIhODYxOTExODU0YTQILCJpYXQiOiE1ODU0ODgyNDAsInNlbnRhdGUiOiJ0eXNpdCI6MTAwMDAwMCB9" :  
        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",  
    }  
}
```

## Results

You can use the tokens to access APIs from other applications, and to execute APIs when the management vCenter Server is down.

# Password Management

# 21

For security reasons, you can change passwords for the accounts that are used by your Cloud Foundation system. Changing these passwords periodically or when certain events occur, such as an administrator leaving your organization, reduces the likelihood of security vulnerabilities.

You specified passwords for your 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 Dashboard. For example:

- Accounts used for service consoles, such as the ESXi root account.
- The single sign-on administrator account.
- The default administrative user account used by virtual appliances.

To provide optimal security and proactively prevent any passwords from expiring, you should rotate passwords every 80 days.

You can also use the VMware Cloud Foundation API to look up and manage credentials. From the SDDC Manager Dashboard, click **Developer Center > API Explorer** and browse to the APIs for managing credentials.

This chapter includes the following topics:

- [Rotate Passwords](#)
- [Manually Update Passwords](#)
- [Look Up Account Credentials](#)
- [Updating SDDC Manager Passwords](#)

## Rotate Passwords

As a security measure, you can rotate passwords for the logical and physical accounts on all racks in your system. The process of password rotation generates randomized passwords for the selected accounts.

You can rotate passwords for the following accounts.

- VxRail Manager
- ESXi

- vCenter Server

By default, the vCenter Server root password expires after 90 days.

- PSC SSO
- NSX-T Edge
- NSX Manager
- vRealize Suite Lifecycle Manager
- SDDC Manager **backup** user

The default password policy for rotated passwords is:

- 15 character 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

To update the SDDC Manager root, super user, and API passwords, see [Updating SDDC Manager Passwords](#).

#### Prerequisites

- Verify that there are no currently failed workflows in your 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 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.

#### Procedure

- 1 From the navigation pane, choose **Administration > Security > Password Management > Locally Managed**.

The Password Management page displays a table with detailed information about all domains, including their account, credential type, FQDN, IP address, and user name. This table is dynamic. Each column can be sorted.

You can click the filter icon next to the table header and filter the results by a string value. For example, click the icon next to **User Name** and enter **admin** to display only domains with that user name value.

- 2 Select the account for which you want to rotate passwords from the **Component** drop-down menu. For example, **ESXI**.

### 3 Select one or more accounts and click **Rotate**.

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. Click on the task name to view sub-tasks. As each of these tasks are run, the status is updated. If the task fails, you can click **Retry**.

#### Results

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.

---

**Note** You can update passwords for **USER** and **SYSTEM** account types.

---

You can update only one password at a time.

Although individual 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: 12
- 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

#### Prerequisites

- Verify that there are no currently failed workflows in your 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 [Chapter 20 User and Group Management](#).

**Procedure**

- 1 From the navigation pane, select **Administration > Security > Password Management**.

The Password Management page displays a table with detailed information about all domains, including their account, credential type, FQDN, IP address, and user name. This table is dynamic. Each column can be sorted.

You can click the filter icon next to the table header and filter the results by a string value. For example, click this icon next to **User Name** and enter **admin** to display only domains with that user name value.

- 2 Select the account whose password you want to update and click **Update** at the top of the page.

---

**Note** If you select more than one account, the **Update** button will be unavailable (dimmed).

---

The Update Password dialog box appears. This dialog box also displays the account name, account type, credential type, and user name, in case you must confirm you have selected the correct account.

- 3 Enter and confirm the new password.

If the passwords do not match, the dialog box displays a red alert.

- 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.

If the Tasks panel shows the task as having failed, click **Retry**.

**Results**

Password updation 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 log in to the SDDC Manager VM using the root account credentials.

**Prerequisites**

Only a user with the ADMIN role can perform this task.

**Procedure**

- 1 SSH in to the SDDC Manager VM using the **vcof** user account.

- 2 (Optional) 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 will be required to enter the user name and the password for a user with the ADMIN role.

- 4 (Optional) 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

You cannot update SDDC Manager passwords through the SDDC Manager Dashboard or by using cURL API requests. Instead, you will need to SSH into the SDDC Manager VM and make the changes there.

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.

- [Update SDDC Manager REST API Account Password](#)

To use the VMware Cloud Foundation API, an API client logs in using the SDDC Manager **admin** account. For security reasons, you should periodically update the password for this account.

- [Update Expired SDDC Manager root Password](#)

This section describes the procedure for updating an expired password for the SDDC Manager root (**root**) user.

## 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 365 days.

### Procedure

- 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:

Option	Description
<b>passwd vcf</b>	To change the super user password.
<b>passwd root</b>	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
```

### Results

The password is updated.

## Update SDDC Manager REST API Account Password

To use the VMware Cloud Foundation API, an API client logs in using the SDDC Manager **admin** account. For security reasons, you should periodically update the password for this account.

If you write a script that invokes the APIs, the script should either prompt the user for the password for the **admin** account or should accept the password as a command line option. As a best practice, you should not encode the password for the account in the script code itself.

Password requirements:

- Length 8-12 characters
- Must include: mix of upper-case and lower-case letters a number a special character such as @ ! # \$ % ^ or ?
- Cannot include: \* { } [ ] ( ) / \ ' " ` ~ , ; : . < >

### Procedure

- 1 SSH in to the SDDC Manager VM using the **vcf** user account.
- 2 Enter **su** to switch to the root user.
- 3 Enter the following command:

```
/opt/vmware/vcf/commonsvcs/scripts/auth/set-basicauth-password.sh admin <password>
```

For *<password>*, enter the new password for the **admin** account.

## Update Expired SDDC Manager root Password

This section describes the procedure for updating an expired password for the SDDC Manager root (**root**) user.

## Prerequisites

### Procedure

- 1 Log into the vSphere Client and select the SDDC Manager VM from Mgmt-ResourcePool.
- 2 From the panel on the right side of the window, select **Summary** and click on **Launch Web Console**.
- 3 In the popup window, select web console and click **OK**. This will open the console in a new browser tab.

The console opens in a new browser window.

- 4 Click **Login**.
- 5 Type **root** as the user name and enter the current password for the root user.
- 6 When prompted for current password, enter the current password.
- 7 When prompted for a new password, enter a different password than the previous one and click **OK**.



# Backing Up and Restoring SDDC Manager and NSX Manager

# 22

It is critically important that you back up the management VMs regularly 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.

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 general procedures on backing up and restoring a full-stack SDDC, please see *VMware Validated Design Backup and Restore*, available from the [VMware Validated Design Documentation](#) page.

---

**Note** This section does not include procedures for updating SDDC Manager state after restoring other Cloud Foundation products. Please contact VMware Support if you need to restore such a product.

---

Follow the best practices below:

- Schedule backups when no other workflows are running.
- Take periodic backups on a daily to weekly frequency.
- If a workflow does not complete successfully and the environment is in this state when the scheduled backup is taken, resolve the failure as soon possible and take an unscheduled backup. Restoring your environment from a backup that includes unresolved failures is more difficult than restoring from a clean backup.

A workflow is resolved when the environment is not in an intermediate state. Some workflows can only be resolved by fixing the failure conditions and retrying the operation. Other workflows can also be resolved by invoking the corresponding delete operation. For example, if adding a host to a workload domain fails, either fix the condition that caused the workflow to fail, or run the workflow that removes the host from the cluster. Contact VMware Support if you are unable to resolve a workflow.

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 using APIs, and are not using composable servers or stretched clusters.

By default, NSX Manager file-based backups are taken on the SFTP server that is built into SDDC Manager.

You should register an external SFTP server with SDDC Manager after you deploy VMware Cloud Foundation for the following reasons:

- An external SFTP server is a prerequisite for using SDDC Manager file-based backups, and you can't enable them until you register a SFTP server.
- By default, NSX Manager file-based backups are taken on the SFTP server that is built into SDDC Manager. Using an external SFTP server provides better protection against failures because it decouples the NSX backups from the SDDC Manager backups. The built-in SFTP server provides temporary protection against failures and should be used while you are setting up an external SFTP server

This chapter includes the following topics:

- [Image-Based Backup and Restore](#)
- [File-Based Backup and Restore](#)

## Image-Based Backup and Restore

For an image-based backup of the SDDC Manager, 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 the backup to a remote site. 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.

For an SDDC Manager backup, connect your backup with the management domain vCenter Server. Configure the product to take non-quieted backups of SDDC Manager. To reduce the backup time and storage cost, use incremental backups in addition to full ones.

## File-Based Backup and Restore

You can use a file-based backup and restore solution for SDDC Manager and NSX Manager.

In a file-based solution, the state of a product is periodically exported to a file that is stored in a different domain than the one where the product is running. If the product needs to be restored, the OVA is redeployed and a selected backup file is used to restore the state. Finally, the post-restore steps are done.

In case you have to restore the SDDC Manager VM or an NSX Manager VM, you select the backup file to restore and download the appropriate OVA file. You can deploy this OVA either through vCenter Server or through the OVF tool. You then load the state on the newly deployed VM.

Note the following limitations for file-based backup:

- This solution requires that you register an external SFTP server. See [Configure an External SFTP Server for File-Based Backups](#). If you do not use an external SFTP server, NSX Managers continue to write backups to the built-in SFTP server on the SDDC Manager VM. This process does not back up the NSX Manager backup files, which leave a gap in protection.
- For the SDDC Manager VM, you must set a backup schedule after setting up an external SFTP server. Backups are not configured automatically. See [Configure a Backup Schedule for SDDC Manager VM](#).
- Neither SDDC Manager nor NSX Manager currently manage the files they back up. It is your responsibility to delete the files that are backed up once their age exceeds your company's retention policy.
- This solution cannot be used for composable servers.
- This solution cannot be used when you have stretched clusters in your environment.

## Configure an External SFTP Server for File-Based Backups

VMware Cloud Foundation allows you to register an external SFTP server with SDDC Manager for backing up NSX Managers and the SDDC Manager VM.

It is important to deploy a reliable SFTP server and ensure it is accessible from the VMware Cloud Foundation instance. If the SFTP server is not available when the SDDC Manager VM or an NSX Manager attempts to back up its state, the backup will not be taken, and any recent changes are not backed up until the retries succeed. To ensure that this situation does not occur, it is recommended that you periodically check that backups are successfully taken, and monitoring that the backups for other products are also being successfully taken. If the SFTP server is not available at the time of deploying a workload domain or upgrading NSX, these operations fail.

When you configure an external SFTP server, SDDC Manager saves the SFTP server details, and then configures the SDDC Manager and all existing NSX Managers to use the SFTP server. When subsequent NSX Managers are deployed, SDDC Manager configures them to use this SFTP server as well.

When you configure an external SFTP server, NSX Manager backups are automatically scheduled at regular intervals. You can check and modify the backup interval in the NSX Manager UI. SDDC Manager VM backups are not scheduled automatically. Use the Cloud Foundation API to set a backup schedule for the SDDC Manager VM. See [Configure a Backup Schedule for SDDC Manager VM](#).

To configure an external SFTP server, perform the following steps:

### Prerequisites

- The external SFTP server must support ECDSA SSH public key.
- Only a user with the ADMIN role can perform this task. See [Chapter 20 User and Group Management](#).

- You will need the SHA256 fingerprint of RSA key of the SFTP server.

#### Procedure

- 1 In the SDDC Manager dashboard, select **Administration > Backup Configuration**.
- 2 Click **+Register External**.
- 3 Enter the IP address of the backup server. Ensure that the server is available for the successful configuration.
- 4 Enter the port number at which the SFTP service is running.
- 5 Enter the credentials of the server.
- 6 Enter the backup directory path of the server. Ensure that the user you specify in step 5 can access the directory path since the backups are saved to this location. It is recommended to provide different directory paths for the different VMware Cloud Foundation instances in case you are using the same SFTP server across all.
- 7 Confirm the fingerprint that is auto populated for the given IP address and the port.
- 8 Enter the pass phrase which is used for both NSX Manager and SDDC Manager backups.
- 9 Click **Save**.
- 10 Click **Confirm**.
- 11 If you have to edit the backup configuration information, perform the following steps:
  - a On the SDDC Manager dashboard, select **Administration > Backup Configuration**.
  - b Click **Edit**.
  - c Edit the text boxes as per your requirement. If there is any change in the IP address or the backup directory path, if you save the configuration, the existing backups are not copied to the new location. Copy them manually.
  - d Enter the backup server password and the passphrase. If there is any change in the passphrase to the existing, you need to use the old passphrase while restoring previously taken backups.
  - e Click **Save**.
  - f Click **Confirm**.

## Configure a Backup Schedule for SDDC Manager VM

Use the Cloud Foundation API to configure a backup schedule for the SDDC Manager VM.

This procedure uses the Cloud Foundation API, which is secured by token-based authentication.

#### Prerequisites

- You must have configured an external SFTP server for file-based backups. See [Configure an External SFTP Server for File-Based Backups](#).

- Only a user with the ADMIN role can perform this task. See [Chapter 20 User and Group Management](#).

## Procedure

- 1 To obtain an access token, run the following command:

```
curl 'https://SDDC_MANAGER_IP/v1/tokens' -k -X POST -H 'Content-Type: application/json' -H 'Accept: application/json' -d '{"username" : "user_name","password" : "user_password"}'
```

Replace the SDDC Manager IP address, user name, and password with the information for your environment.

The command returns an access token and a refresh token.

```
{
  "accessToken": "eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiI2ZGRmZmU3MC0zZGEwLTQ5NmU0OTQzMCIhOGNkNTQ0YT  
k2ZGMIkCjYXQiojElODU3ODE3ODcsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbCIsIm  
lzc2cyI6InZjZ  
ilhdXRoIiwiaXVkiIjoic2RkYy1zZXJ2aWNlcyIsIm5iZiI6MTU0NTc4MTC4NywiZXBwIjoxNTg1NzglMzg3LCJ1c2Vy  
IjoiaWRtaW5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbCI  
sInNjb3BlIjpbIkJBQ0tVUF9DT05GSUdfUkVBRClIsIkNSRURFTlRJQUxfUkVBRClSlVTRVJfVlJJVEUiLCJPEhF  
UL9XUklURSI6IkJBQ0tVUF9DT05GSUdfVlJJVEUiLCJPEhFUL9SRUFEIiwiaWF0Ij05OTQzMCIhOGNkNTQ0YT  
klURSIJdfQ.WCpUPRIm5A6X_406HTJF7TbTSa0g9l_AQbt7OcBPb1m", "refreshToken":
  {
    "id": "47c07f35-0a89-4df5-a3a3-f31265ebbb7a"
  }
}
```

- 2** Set the backup schedule, for example:

[illegible]

Option	Description
Resource Type	Enter <b>SDDC_MANAGER</b> .
Frequency	Enter <b>HOURLY</b> or <b>WEEKLY</b> .
Days of Week	<p>Enter the days to back up SDDC Manager. When selecting multiple days, use a comma to separate them. For example: "daysOfWeek" : [ "SUNDAY", "THURSDAY" ]</p> <p><b>Note</b> Only available if the frequency is set to <b>WEEKLY</b>.</p>
Hour of Day	Enter the hour of day to perform the backup.
Minute of Hour	Enter the minute of the hour to perform the backup.

Note the `<id>` that gets returned.

### 3 Track the status of the backup schedule configuration.

```
curl -H "Authorization: Bearer eyJhbGciOiJIUzI1NiJ9.eyJqdGkiOiIxNzMyYThmOC00NGQwLTQ5MzYtYjJqWmC0xMzc5NzMyMjdmOWUiLCJpYXQiOiE1ODUzODQ5MTgsInN1YiI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbnBhcisImlscyT6InZjZilhdXRoIiwiaXYXVkaioic2RkYy1zZXJ2aWNlcysIm5iZiI6MTU4NTc4NDkxOCwiZXhwIjoxtNg1Nzg4NTE4LClc1c2VyIjoiYWRTaw5pc3RyYXRvckB2c3BoZXJlLmxvY2FsIiwibmFtZSI6ImFkbWluaXN0cmF0b3JAdnNwaGVyZS5sb2NhbnBhcisInnjb3BlIjpjbIkjBQ0tVUF9DT05GSudfUkVBRCIsIkNSRURFTlRJQUxfukVBRCIsIlVTRVFjfv1JJVEUilCJPVEhfU19XUKlURSIIsIkjBQ0tVUF9DT05GSudfv1JJVEUilCJPVEhfU19SRUFEIiwivVNfU19SRUFEIiwqIjFREVOVElBTf9XUklURSjdFQ.Ya4XSzntsRHUZFRBNKGY7Js6xrGYGe8KdgJ2QbihFmg" -H 'Content-Type: application/json' -k -X GET https://SDCC_MANAGER_IP/v1/tasks/<id>|json_pp
```

Replace `<id>` with the ID from the previous step.

4 Repeat the previous step until the task status is `SUCCESSFUL`.

## Restore SDDC Manager

See the Backup and Restore Section of the *VMware Cloud Foundation API Reference Guide* for the manual procedure to restore SDDC Manager from file-based backups using the Cloud Foundation APIs.

# Cloud Foundation Glossary

# 23

Term	Description
availability zone	Collection of infrastructure components. Each availability zone is isolated from other availability zones to prevent the propagation of failure or outage across the data center.
Application virtual networks (AVNs)	Virtual networks backed by overlay segments using the encapsulation protocol of NSX-T. Virtual Networks use a single IP network address space, to span across data centers.
bring-up	Initial configuration of a newly deployed Cloud Foundation system. During the bring-up process, the management domain is created and the Cloud Foundation software stack is deployed on the management domain.
cluster image	Precise description of the software, components, vendor add-ons, and firmware to run on a host. With this new functionality, you set up a single image and apply it to all hosts in a cluster, thus ensuring cluster-wide host image homogeneity.
commission host	Adding a host to Cloud Foundation inventory. The host remains in the free pool until it is assigned to a workload domain.
composability	Ability to dynamically configure servers to meet the needs of your workloads without physically moving any hardware components. You bind disaggregated hardware components (compute, network, storage, and offload components) together to create a logical system based on the needs of your applications.
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 cleaned up.
decommission host	Remove an unassigned host from the Cloud Foundation inventory. SDDC Manager does not manage decommissioned hosts.
Edge cluster	A logical grouping of Edge nodes. These nodes run on a vSphere cluster, and provide north-south routing and network services for the management and VI workload domains.
free pool	Hosts in the Cloud Foundation inventory that are not assigned to a workload domain
host	An imaged server.
inventory	Logical and physical entities managed by Cloud Foundation.
Kubernetes - Workload Management	With Kubernetes - Workload Management, you can deploy and operate the compute, networking, and storage infrastructure for vSphere with Kubernetes workloads. A vSphere with Kubernetes 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.

Term	Description
management domain	Cluster of physical hosts that contains the management component VMs
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.
patch update bundle	Contains bits to update the appropriate Cloud Foundation software components in your management or VI workload domain.
region	A Cloud Foundation instance.
SDDC Manager	Software component that provisions, manages, and monitors the logical and physical resources of a Cloud Foundation system.
SDDC Manager VM	Virtual machine (VM) that contains the SDDC Manager services and a shell from which command line tools can be run. This VM exposes the SDDC Manager UI.
server	Bare metal server in a physical rack. After imaging, it is referred to as a host.
unassigned host	Host in the free pool that does not belong to a workload domain.
vSphere Lifecycle Manager (vLCM)	A vCenter service, which is now integrated with Cloud Foundation, that enables centralized and simplified lifecycle management of ESXi hosts.
workload domain	A policy based resource container with specific availability and performance attributes that combines vSphere, storage (vSAN, NFS, or VMFS on FC) and networking (NSX-T) into a single consumable entity. A workload domain can be created, expanded, and deleted as part of the SDDC lifecycle operations. It can contain cluster(s) of physical hosts with a corresponding vCenter to manage them. The vCenter for a workload domain physically lives in the management domain.