

# VMware Cloud Director Tenant Portal Guide

Modified on 4 APR 2021  
VMware Cloud Director 10.2

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 © 2017-2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

VMware Cloud Director™ Tenant Portal Guide 10

## 1 Getting Started with the VMware Cloud Director Tenant Portal 12

- Understanding VMware Cloud Director™ 12
- Log In to the VMware Cloud Director Tenant Portal 14
- VMware Cloud Director Tenant Portal Roles and Rights 14
- Using the VMware Cloud Director Tenant Portal 15
- Use the VMware Cloud Director Global Search 16
- Use the VMware Cloud Director Quick Search 17
- View Tasks 18
- Stop a Task in Progress 19
- View Events 19
- Set User Preferences 20

## 2 Working with Virtual Machines 21

- Virtual Machine Architecture 22
- Virtual Machine Encryption 23
- View Virtual Machines 24
- Create a New Standalone Virtual Machine 25
- Fast Provisioning of Virtual Machines 26
- Opening a Virtual Machine Console 27
  - Install VMware Remote Console on a Client 27
  - Open a Virtual Machine Remote Console 28
  - Open a Web Console 28
- Performing Power Operations on Virtual Machines 29
  - Power On a Virtual Machine 29
  - Power Off a Virtual Machine 30
  - Shut Down a Guest Operating System 30
  - Reset a Virtual Machine 31
  - Suspend a Virtual Machine 31
  - Discard the Suspended State of a Virtual Machine 32
  - Power On Multiple VMs 32
  - Power Off Multiple Virtual Machines 32
  - Reset Multiple Virtual Machines 33
- Install VMware Tools in a Virtual Machine 33
- Upgrade the Virtual Hardware Version for a Virtual Machine 34
- Edit Virtual Machine Properties 35
  - Change the General Properties of a Virtual Machine 35

- Change the Hardware Properties of a Virtual Machine 36
- Change the Guest OS Customization Properties of a Virtual Machine 39
- Change the Advanced Properties of a Virtual Machine 43
- Insert Media 45
- Eject Media 46
- Copy a Virtual Machine to a Different vApp 46
- Move a Virtual Machine to a Different vApp 47
- Virtual Machine Affinity and Anti-Affinity 48
  - View Affinity and Anti-Affinity Rules 48
  - Create an Affinity Rule 48
  - Create an Anti-Affinity Rule 49
  - Edit an Affinity or Anti-Affinity Rule 50
  - Delete an Affinity or Anti-Affinity Rule 50
- Monitor Virtual Machines 50
- Working with Snapshots 52
  - Take a Snapshot of a Virtual Machine 52
  - Revert a Virtual Machine to a Snapshot 53
  - Remove a Snapshot of a Virtual Machine 54
- Renew a Virtual Machine Lease 54
- Delete a Virtual Machine 54
- Auto Scale Groups 55
  - Create a Scale Group 55
  - Add an Auto Scaling Rule 56

### 3 Working with vApps 58

- View vApps 59
- Build a New vApp 59
- Create a vApp From an OVF Package 61
- Add a vApp from a Catalog 64
- Create a vApp from a vApp Template 65
- Import a Virtual Machine from vCenter Server as a vApp 67
- Performing Power Operations on vApps 67
  - Power on a vApp 67
  - Power off a vApp 68
  - Reset a vApp 68
  - Suspend a vApp 69
  - Discard the Suspended State of a vApp 69
  - Power on Multiple vApps 70
  - Power off Multiple vApps 70
  - Discard the Suspended State of Multiple vApps 70
  - Reset Multiple vApps 71

Suspend Multiple vApps	71
Open a vApp	72
Edit vApp Properties	72
Edit the General Properties of the vApp	72
Edit the Start and Stop Order of Virtual Machines in a vApp	73
Edit the Guest Properties of a vApp	74
Share a vApp	75
Display a vApp Network Diagram	75
Working with Networks in a vApp	76
View vApp Networks	77
Fence a vApp Network	77
Add a Network to a vApp	78
Configuring Network Services for a vApp Network	79
Delete a vApp Network	86
Working with Snapshots	86
Take a Snapshot of a vApp	86
Revert a vApp to a Snapshot	87
Remove a Snapshot of a vApp	88
Take Snapshots of Multiple vApps	88
Remove the Snapshots of Multiple vApps	89
Revert Multiple vApps to Snapshots	89
Change the Owner of a vApp	90
Move a vApp to Another Virtual Data Center	90
Copy a Stopped vApp to Another Virtual Data Center	91
Copy a Powered-On vApp	92
Add a Virtual Machine to a vApp	92
Save a vApp as a vApp Template to a Catalog	93
Download a vApp as an OVF Package	95
Renew a vApp Lease	95
Delete a vApp	96
Delete Multiple vApps	96
<b>4 Working with Kubernetes Clusters</b>	<b>98</b>
Add an Organization VDC Kubernetes Policy	99
Edit an Organization VDC Kubernetes Policy	101
Create a Tanzu Kubernetes Cluster	101
Create a Native Kubernetes Cluster	103
Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster	105
Configure External Access to a Service in a Tanzu Kubernetes Cluster	106
<b>5 Working with Networks</b>	<b>108</b>

Managing Organization Virtual Data Center Networks	111
View the Available Organization VDC Networks	112
Add an Isolated Organization Virtual Data Center Network	112
Add a Routed Organization Virtual Data Center Network	114
Add a Direct Organization Virtual Data Center Network	116
Add an Organization VDC Network with an Imported NSX-T Data Center Logical Switch	117
Edit the General Settings of an Organization Virtual Data Center Network	118
Connect an Organization Virtual Data Center Network to an Edge Gateway	118
Disconnect an Organization VCD Network from an Edge Gateway	119
Convert the Interface of a Routed Organization VDC Network	120
View the IP Addresses Used for an Organization Virtual Data Center Network	120
Add IP Addresses to an Organization Virtual Data Center Network IP Pool	121
Edit or Remove IP Ranges Used in an Organization Virtual Data Center Network	121
Edit the DNS Settings of an Organization Virtual Data Center Network	122
Configure DHCP Settings for an Isolated Organization Virtual Data Center Network	122
Add a DHCP Pool to a Routed Organization Virtual Data Center Network Backed by NSX-T Data Center	123
Edit or Delete an Existing DHCP Pool for an Isolated Organization Virtual Data Center Network Backed by NSX Data Center for vSphere	124
Reset an Organization Virtual Data Center Network	125
Delete an Organization Virtual Data Center Network	125
Managing Data Center Group Networking with NSX-T Data Center	126
Managing Data Center Groups with an NSX-T Data Center Network Provider Type	127
Using Distributed Firewall in a Data Center Group with an NSX-T Data Center Network Provider Type	128
Managing Data Center Group Networks with an NSX-T Data Center Network Provider Type	133
Managing Egress Points for Data Center Groups with an NSX-T Data Center Network Provider Type	138
Managing Data Center Group Networking with NSX Data Center for vSphere	140
Managing Data Center Groups with NSX Data Center for vSphere Network Provider Type	141
Managing Data Center Group Networks Backed by NSX Data Center for vSphere	155
Managing NSX Data Center for vSphere Edge Gateway Services	157
Getting Started with VMware Cloud Director Advanced Networking with NSX Data Center for vSphere	158
Tenant Firewall Configuration with NSX Data Center for vSphere	158
Managing NSX Data Center for vSphere Edge Gateway DHCP	169
Managing Network Address Translation on an NSX Data Center for vSphere Edge Gateway	174
Advanced Routing Configuration for NSX Data Center for vSphere Edge Gateways	177
Load Balancing with NSX Data Center for vSphere	186
Configure Secure Access Using VPN on an NSX Data Center for vSphere Edge Gateway	198

SSL Certificate Management on an NSX Data Center for vSphere Edge Gateway	224
Custom Grouping Objects for NSX Data Center for vSphere Edge Gateways	230
Statistics and Logs for an NSX Data Center for vSphere Edge Gateway	234
Enable SSH Command-Line Access to an NSX Data Center for vSphere Edge Gateway	236
Working with Security Tags for NSX Data Center for vSphere Edge Gateways	236
Working with Security Groups for NSX Data Center for vSphere Edge Gateways	240
Managing NSX-T Data Center Edge Gateways	244
Add an IP Set to an NSX-T Data Center Edge Gateway	244
Add an NSX-T Data Center Edge Gateway Firewall Rule	245
Add an SNAT or a DNAT Rule to an NSX-T Edge Gateway	246
Configure a DNS Forwarder Service on an NSX-T Edge Gateway	249
Create Custom Application Port Profiles	249
IPsec Policy-Based VPN for NSX-T Data Center Edge Gateways	250
Configure Dedicated External Network Services	253
Working with NSX Advanced Load Balancing	258
<b>6 Using Named Disks and Reviewing Storage Policies</b>	<b>265</b>
Creating and Using Named Disks	265
Create a Named Disk	265
Edit a Named Disk	266
Attach a Named Disk to a Virtual Machine	267
Delete a Named Disk	267
Review Storage Policy Properties	267
<b>7 Reviewing and Editing Virtual Data Center Properties</b>	<b>269</b>
Review Virtual Data Center Properties	269
Review the Virtual Data Center Metadata	269
Limit Access to an Organization VDC to Specific Users and Groups in Your Organization	270
<b>8 Working with Dedicated vCenter Server Instances, Endpoints, and Proxies</b>	<b>272</b>
Using Chrome Browser Extension for VMware Cloud Director	273
Configure Your Browser with Your Proxy Settings	273
Log In to the UI of a Component by Using an Endpoint	274
<b>9 Working with vApp Templates</b>	<b>275</b>
View a vApp Template	275
Create a vApp Template from an OVF File	276
Import a Virtual Machine from vCenter Server as a vApp Template	277
Assign a VM Placement Policy and a VM Sizing Policy to a vApp Template	277
Download a vApp Template	278
Delete a vApp Template	279

## 10 Working with Media Files 280

Upload Media Files 280

Delete a Media File 281

Download a Media File 281

## 11 Working with Catalogs 282

View Catalogs 283

Create a Catalog 283

Share a Catalog 284

Delete a Catalog 285

Change the Owner of a Catalog 285

Manage Metadata for a Catalog 286

Publish a Catalog 287

Subscribe to an External Catalog 287

Update the Location URL and the Password for a Subscribed Catalog 288

Synchronize a Subscribed Catalog 288

## 12 Working with Organization Virtual Data Center Templates 290

View Available Virtual Data Center Templates 290

Instantiate a Virtual Data Center from a Template 291

## 13 Managing Users, Groups and Roles 292

Managing Users 292

Create a User 292

Import Users 294

Modify a User 295

Deactivate or Activate a User Account 295

Delete a User 296

Unlock a Locked Out User Account 296

Manage the Resource Quotas of a User 296

Managing Groups 297

Import a Group 297

Delete a Group 298

Edit a Group 299

Manage the Resource Quotas of a Group 299

Roles and Rights 300

Predefined Roles and Their Rights 300

Rights in Predefined Global Tenant Roles 302

Create a Custom Tenant Role 307

Edit a Custom Tenant Role 308

Delete a Role 309

## **14** Configure Identity Providers 310

Enable Your Organization to Use a SAML Identity Provider 310

Edit LDAP Settings for Your Organization 312

Configure, Test, and Synchronize an LDAP Connection 313

## **15** Managing Certificates 315

Import Trusted Certificates 315

Import Certificates to the Certificates Library 316

## **16** Managing Your Organization 317

Edit the Organization Name and Description 317

Modify Your Email Settings 318

Test SMTP Settings 319

Modify Domain Settings for the Virtual Machines in Your Organization 319

Working with Multiple Sites 320

Configure and Manage Multisite Deployments 320

Understanding Leases 321

Modify the vApp and vApp Template Lease Policies Within Your Organization 322

Modify the Password and User Account Policies Within Your Organization 323

Create an Advisories Dashboard 323

## **17** Working with Service Library 325

Search for a Service 325

Execute a Service 326

## **18** Managing Defined Entities 327

Working with Custom Entity Definitions 329

Search for a Custom Entity 329

Edit a Custom Entity Definition 330

Add a Custom Entity Definition 330

Custom Entity Instances 331

Associate an Action to a Custom Entity 332

Dissociate an Action from a Custom Entity Definition 332

Publish a Custom Entity 333

Delete a Custom Entity 333

# VMware Cloud Director™ Tenant Portal Guide

The *VMware Cloud Director™ Tenant Portal Guide* provides information about how to use the VMware Cloud Director tenant portal. In this release, you use the tenant portal to administrate your organization, create and configure virtual machines, vApps, and networks within vApps. You can also configure advanced networking capabilities that are provided by VMware NSX® for vSphere® within a VMware Cloud Director environment. With the VMware Cloud Director tenant portal, you can also create and manage catalogs, vApp and VDC templates, and create and manage cross-virtual data center networks.

## Intended Audience

This guide is intended for anyone who wants to use the capabilities of the VMware Cloud Director tenant portal. The information is written primarily for **organization administrators** who use the tenant portal to administer their organization, manage virtual machines, vApps, networks, and so on.

## VMware Technical Publications Glossary

VMware Technical Publications provides a glossary of terms that might be unfamiliar to you. For definitions of terms as they are used in VMware technical documentation, go to <http://www.vmware.com/support/pubs>.

## Usage Terms & Conditions

VMware gives you permission to modify this tenant user guide (the “Guide”) as reasonably needed to customize it to reflect your operational processes, and then to reproduce and distribute the modified Guide to your customers. You may not charge your customers a fee for access to the modified Guide. YOU ACKNOWLEDGE THAT THE GUIDE IS PROVIDED TO YOU WITHOUT CHARGE, “AS IS” WITHOUT WARRANTY OF ANY KIND, AND ONLY FOR THE PURPOSE DESCRIBED ABOVE. ACCORDINGLY, THE TOTAL LIABILITY OF VMWARE AND ITS SUPPLIERS ARISING OUT OF OR RELATED TO PROVIDING YOU WITH ACCESS TO THE GUIDE SHALL NOT EXCEED \$100. IN NO EVENT SHALL VMWARE OR ITS SUPPLIERS HAVE LIABILITY FOR ANY INDIRECT, INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS

INTERRUPTION, OR LOSS OF BUSINESS INFORMATION), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, EVEN IF VMWARE OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. THESE LIMITATIONS SHALL APPLY NOTWITHSTANDING ANY FAILURE OF ESSENTIAL PURPOSE OF ANY LIMITED REMEDY.

# Getting Started with the VMware Cloud Director Tenant Portal

# 1

When you log into the tenant portal, there are a number of tasks you can complete, from creating virtual machines and vApps, to setting up advanced networking configuration and running vRealize Orchestrator workflows.

This chapter includes the following topics:

- [Understanding VMware Cloud Director™](#)
- [Log In to the VMware Cloud Director Tenant Portal](#)
- [VMware Cloud Director Tenant Portal Roles and Rights](#)
- [Using the VMware Cloud Director Tenant Portal](#)
- [Use the VMware Cloud Director Global Search](#)
- [Use the VMware Cloud Director Quick Search](#)
- [View Tasks](#)
- [Stop a Task in Progress](#)
- [View Events](#)
- [Set User Preferences](#)

## Understanding VMware Cloud Director™

VMware Cloud Director™ provides role-based access to a web-based tenant portal that allows the members of an organization to interact with the organization's resources to create and work with vApps and virtual machines.

Before you can access your organization, a VMware Cloud Director **system administrator** must create the organization, assign it resources, and provide the URL to access the tenant portal. Each organization includes one or more **organization administrators**, who finish setting up the organization by adding members and setting policies and preferences. After the organization is set up, non-administrator users can log in to create, use, and manage virtual machines and vApps.

## Organizations

An organization is a unit of administration for a collection of users, groups, and computing resources. Users authenticate at the organization level, supplying credentials established by an **organization administrator** when the user was created or imported. **System administrators** create and provision organizations, while **organization administrators** manage organization users, groups, and catalogs.

## Users and Groups

An organization can contain an arbitrary number of users and groups. Users can be created locally by the organization administrator or imported from a directory service. Groups must be imported from the directory service. Permissions within an organization are controlled through the assignment of rights and roles to users and groups.

## Virtual Data Centers

An organization virtual data center provides resources to an organization. Virtual data centers provide an environment where virtual systems can be stored, deployed, and operated. They also provide storage for virtual CD and DVD media. An organization can have multiple virtual data centers.

## Organization Virtual Data Center Networks

An organization virtual data center network is contained within a VMware Cloud Director organization virtual data center and is available to all the vApps in the organization. An organization virtual data center network allows vApps within an organization to communicate with each other. An organization virtual data center network can be connected to an external network or isolated and internal to the organization. Only **system administrators** can create organization virtual data center networks, but **organization administrators** can manage organization virtual data center networks, including the network services they provide.

## vApp Networks

A vApp network is contained within a vApp and allows virtual machines in the vApp to communicate with each other. You can connect a vApp network to an organization virtual data center network to allow the vApp to communicate with other vApps in the organization and outside of the organization, if the organization virtual data center network is connected to an external network.

## Catalogs

Organizations use catalogs to store vApp templates and media files. The members of an organization that have access to a catalog can use its vApp templates and media files to create their own vApps. **Organization administrators** can copy items from public catalogs to their organization catalog.

## Dedicated vCenter Server Instances (SDDCs) and Proxies

A Software-Defined Data Center (SDDC) encapsulates an entire vCenter Server environment. A dedicated vCenter Server instance can include one or more proxies that provide access to different components from the underlying environment. The **system administrator** can publish one or more dedicated vCenter Server instances to your organization. You can use the containing proxies to access the UI or API of the proxied components.

## Log In to the VMware Cloud Director Tenant Portal

You can access the VMware Cloud Director tenant portal by using a URL that is specific to your organization.

Contact your **organization administrator** if you do not know the tenant portal organization URL. See the *VMware Cloud Director Release Notes* for information about supported browsers and configurations.

### Procedure

- 1 In a Web browser, navigate to the tenant portal URL of your organization.  
For example, *https://cloud.example.com/tenant/myOrg*.
- 2 Enter your user name and password, and click **Log In**.

## VMware Cloud Director Tenant Portal Roles and Rights

VMware Cloud Director includes a preconfigured set of user roles and their rights. The roles that are able to access the VMware Cloud Director tenant portal are the roles created by default in any organization or other roles that are created by the organization administrator.

Users who are assigned the following organization roles can access the tenant portal. The items they see and the actions they can perform depend on the rights that are associated with a particular role.

- **Organization Administrator**
- **Catalog Author**
- **vApp Author**
- **vApp User**
- **Console Access Only**

For information about the predefined roles and their rights, see [Predefined Roles and Their Rights](#).

## Using the VMware Cloud Director Tenant Portal

If you have more than one virtual data center, when you log in to the VMware Cloud Director tenant portal, you are navigated to the **Data Centers** dashboard screen. If you have only one virtual data center, when you log in to the VMware Cloud Director tenant portal, you are directly navigated to the data center.

The **Data Centers** dashboard screen is part of the VMware Cloud Director multisite feature that makes it possible for tenants to see their geographically distributed cloud environment as a single entity. For more information about multisite, see [Working with Multiple Sites](#).

The dashboard is a unified view of the VMware Cloud Director virtual data centers and sites not only in a single organization. In a multi-cell and multi-organization environment, you can also see the virtual data centers for all other associated organizations.

---

**Note** Depending on their rights, tenant users can see all member sites of an organization or only a subset of sites.

---

The information about the organization is displayed on top in the summary ribbon.

If you log in as an **organization administrator**, you can see:

- The number of sites, organizations, and virtual data centers
- Total number of running vApps and virtual machines
- Used hardware resources, such as CPU, memory, and storage

The virtual data centers display in a card view. Each card contains information about the organization to which the virtual center belongs, the number of vApps, the total number of virtual machines and the number of virtual machines that are in a running state. The card also shows the available CPU, memory, and storage capacity for the data center and displays real-time metrics about the current allocations and reservations of resources.

From the top navigation, you can navigate to the different menu items.

Menu Item	Description
Data Centers	Navigates you to the <b>Virtual Data Center</b> , <b>Data Center Groups</b> and <b>Dedicated vSphere Datacenters</b> resources in your organization
Virtual Data Center	Navigates you to the <b>Virtual Data Center</b> screen that displays the virtual data centers within the organization.
Dedicated vSphere Data Centers	Navigates you to the screen that displays the dedicated vSphere data centers that your service provider has published to your organization.
Applications	Navigates you to the <b>Virtual Applications</b> and <b>Virtual Machines</b> resources in your organization.
Libraries	Navigates you to a consolidated view for vApp templates, catalogs, media, and other types of files. You use these templates and files to deploy virtual machines or vApps.
Networking	Navigates you to the networks, edge gateways and data center groups in your organization.

Menu Item	Description
Administration	Navigates you to the <b>Access Control</b> , <b>Identity Provider</b> configuration screens, and to the general, email, guest personalization, metadata, multisite, and policies settings for your organization.
Monitor	Navigates you to the <b>Tasks</b> and <b>Events</b> screens. The <b>Tasks</b> screen displays the tasks reported by VMware Cloud Director. The <b>Events</b> screen that displays the events reported by VMware Cloud Director.

You can customize your VMware Cloud Director Tenant Portal by using the `Branding Cloud Director OpenAPIs`. For information about using the Cloud Director OpenAPI, see the *Getting Started Cloud Director OpenAPI* document at <https://code.vmware.com>.

## Use the VMware Cloud Director Global Search

You can use the VMware Cloud Director global search to perform a search by a name or part of a name among the names of the objects in your environment. You can also search for a virtual machine by its IP address if the IP address of the virtual machine is static.

The list of preset objects is:

- Data centers
- vApp templates
- vApps
- Virtual machines
- vApp networks
- Catalogs

If a virtual machine uses an IP address assigned by DHCP, the search does not return its IP address. If you want to search for a virtual machine that is with an IP address assigned by DHCP, you must search by name.

By default, you can search only within the objects in your local site. If you have a multisite environment, you can search among multiple sites.

### Procedure

- 1 In the right-upper corner of the VMware Cloud Director tenant portal, click the **Search** icon.
- 2 (Optional) Pin the search panel by clicking the **Pin** icon.
- 3 In the **Search** text box, enter a symbol, a part of a name, or IP address by which to search for matching object names or static IP addresses of virtual machines.
- 4 If you use a multisite environment, select the sites within which you want to perform the search.
- 5 Press **Enter**.

## Results

The top five matching results per object type are displayed. The results are sorted alphabetically.

### What to do next

- To see more results, if there are any, click **Load more** under each object type.
- To see more information about a specific object from the search results, point to the object.
- To manage a specific object, for example, to view or modify the settings of an object, click the object. The details about the object display on the left.

## Use the VMware Cloud Director Quick Search

You can use the VMware Cloud Director quick search to find screens, entities, and actions. The results depend on your location in the UI.

The results depend on the context, whether you selected an entity, and depending on the available actions for a particular entity. The search results are grouped into sections.

- **Global Navigation** - the results in this section are not related to a specific entity, for example, Edge Gateways, LDAP, Tasks, Trusted Certificates, Virtual Machines, and so on. You get these results regardless of where you are in the UI.
- **Contextual Navigation** - the results in this section depend on the selected entity in the UI. For example, vApp specific views like VMs, Network Diagram, and so on. If you select an entity like a vApp, the search shows both global and contextual navigation results and any actions that might be applicable to the entity.
- **Contextual Actions** - the results in this section depend on the selected entity in the UI. Depending on your location in the UI and the entity you select, by using the quick search results, you can perform an action related to the entity. For example, searching from the details view of a virtual machine displays results from the global views, contextual views, and actions that you can perform on the selected VM.
- **Entity Search by Name** - if you are viewing a list of entities, the search results can include also names of entities of the same type as the ones in the list. For example, if you are viewing a list of VMs, the search results include global navigation matches and matching names of VMs. If there is more than one page of entities in the list you are viewing, the search checks the full list of entities and might show a name that is not visible on the current page.

### Procedure

- 1 Open the **Quick Search** window.
  - From the top navigation bar, click the **Help** menu and select **Quick Search**.
  - Press Ctrl+. or Cmd+., depending on your operating system.
- 2 Enter search criteria.

- 3 Browse through the results and select an option or perform an action by clicking or pressing Enter.

You can use the up and down arrow keys to browse through the search results.

## View Tasks

From the tenant portal, you can view the list of recent tasks, together with their details and status. In addition, you can also see the list of all tasks.

By default, the **Recent Tasks** panel is displayed at the bottom of the tenant portal and contains a list of the tasks that have been recently run. When you start an operation, for example to create a virtual machine, the task is displayed in the panel. In case you minimize the **Recent Tasks** panel, you still see the number of running or failed recent tasks. You can always open the **Recent Tasks** panel again by clicking the double arrows.

The tasks view lists all tasks and shows when tasks were run, and whether they were successfully completed. This view is the first step for troubleshooting problems in your environment. The tasks view contains long running operations, such as virtual machine or vApp creation.

### Procedure

- 1 In the top navigation bar, click **Monitor** and **Tasks**.

The list of all tasks displays, together with the time the task was run, and the status of the task.

- 2 Click the editor icon (  ) to change the details you want to view about the tasks.
- 3 (Optional) To view the task details, click the name of the task.

The task details include information such as the reason for the failure, when the task has failed, and so on.

Detail	Description
Operation	Name of the performed operation.
Job ID	ID of the task.
Type	The object on which the task was performed. For example, if you created a virtual machine, the type is <code>vm</code> .
Organization	Organization name.
Status	Status of the task, such as Succeeded, Running, or Failed.
Initiator	User who started the operation.
Start time	Date and time when the operation started.
Completion time	Date and time when the operation succeeded or failed.

Detail	Description
Service namespace	Service name, such as <i>com.vmware.cloud</i> .
Details	Reason for the failure of the task. For example, if you try to create a snapshot of a virtual machine, and the operation fails, because the storage is insufficient, the task details are of the type: <i>The requested operation will exceed the VDC's storage quota: storage policy "*" has 8,693 MB remaining, requested 41,472 MB.</i>

## Stop a Task in Progress

If you accidentally start an operation before applying or reviewing all necessary settings, you can stop the task in progress.

By default, the **Recent Tasks** panel is displayed at the bottom of the portal. When you start an operation, for example to create a virtual machine, the task is displayed in the panel.

### Prerequisites

The **Recent Tasks** panel must be open.

### Procedure

- 1 Start a long-running operation.

Long-running operations are operations such as creating a virtual machine or a vApp, power operations performed on virtual machines and vApps, and so on.

- 2 In the **Recent Tasks** panel, click the **Cancel** icon.
- 3 In the **Cancel Task** dialog box, confirm that you want to cancel the task by clicking **OK**.

### Results

The operation stops.

## View Events

From the portal, you can view the list of all events, as well as their details and status.

The events view is a way to view the status of the events in your portal. The view shows when the events happened, and whether they were successful. The events view contains one-time occurrences, such as user logins and object creation, or deletion.

### Procedure

- 1 In the top navigation bar, click **Monitor** and **Events**.

The list of all events displays, along with the time the event happened and the status of the event.

- 2 Click the editor icon (  ) to change the details you want to view about the events.

### 3 (Optional) Click an event to view the event details.

Detail	Description
Event	Name of the event. For example, if you modify a vApp to include virtual machines in it, the event that starts the whole operation is <i>Task 'Modify vApp' start</i> .
Event ID	ID of the task.
Type	The object on which the task was performed. For example, if you created a virtual machine, the type is <i>vm</i> .
Target	Target object of the event. For example, when you modify a vApp to include virtual machines in it, the target of the <i>Task 'Modify vApp' start</i> event is <i>vdcUpdateVapp</i> .
Status	Status of the event, such as Succeeded or Failed.
Service namespace	Service name, such as <i>com.vmware.cloud</i> .
Organization	Name of the organization.
Owner	User who triggered the event.
Time of occurrence	Date and time when the event occurred.

## Set User Preferences

You can set certain display and system alert preferences that take effect every time you log in to the system.

To learn more about leases, see [Understanding Leases](#).

### Procedure

- 1 In the top navigation bar, click your user name and select **User preferences**.
- 2 Select the page to appear when you log in.
  - a Select the radio button next to **Start Page** and click **Edit**.
  - b Select an option from the drop-down menu and click **Save**.
- 3 Configure an email notification for runtime lease expirations.
  - a Select the radio button next to **Deployment Lease Alert Time** and click **Edit**.
  - b Enter a value in seconds and click **Save**.
- 4 Configure an email notification for storage lease expirations.
  - a Select the radio button next to **Storage Lease Alert Time** and click **Edit**.
  - b Enter a value in seconds and click **Save**.

# Working with Virtual Machines

# 2

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. The virtual machine consists of a set of specification and configuration files, and is backed by the physical resources of a host. Every virtual machine has virtual devices that provide the same functionality as physical hardware but are more portable, more secure, and easier to manage.

In addition to the operations that you can run on a physical machine, VMware Cloud Director virtual machines support virtual infrastructure operations, such as taking a snapshot of virtual machine state, and moving a virtual machine from one host to another.

Starting with VMware Cloud Director 9.5, virtual machines support IPv6 connectivity. You can assign IPv6 addresses to virtual machines connected to IPv6 networks.

---

**Important** All steps for working with virtual machines are documented from the card view, assuming that you have more than one virtual data center. Completing the same procedures from the grid view is also possible, but the steps might slightly vary.

---

This chapter includes the following topics:

- [Virtual Machine Architecture](#)
- [Virtual Machine Encryption](#)
- [View Virtual Machines](#)
- [Create a New Standalone Virtual Machine](#)
- [Fast Provisioning of Virtual Machines](#)
- [Opening a Virtual Machine Console](#)
- [Performing Power Operations on Virtual Machines](#)
- [Install VMware Tools in a Virtual Machine](#)
- [Upgrade the Virtual Hardware Version for a Virtual Machine](#)
- [Edit Virtual Machine Properties](#)
- [Insert Media](#)
- [Eject Media](#)
- [Copy a Virtual Machine to a Different vApp](#)

- [Move a Virtual Machine to a Different vApp](#)
- [Virtual Machine Affinity and Anti-Affinity](#)
- [Monitor Virtual Machines](#)
- [Working with Snapshots](#)
- [Renew a Virtual Machine Lease](#)
- [Delete a Virtual Machine](#)
- [Auto Scale Groups](#)

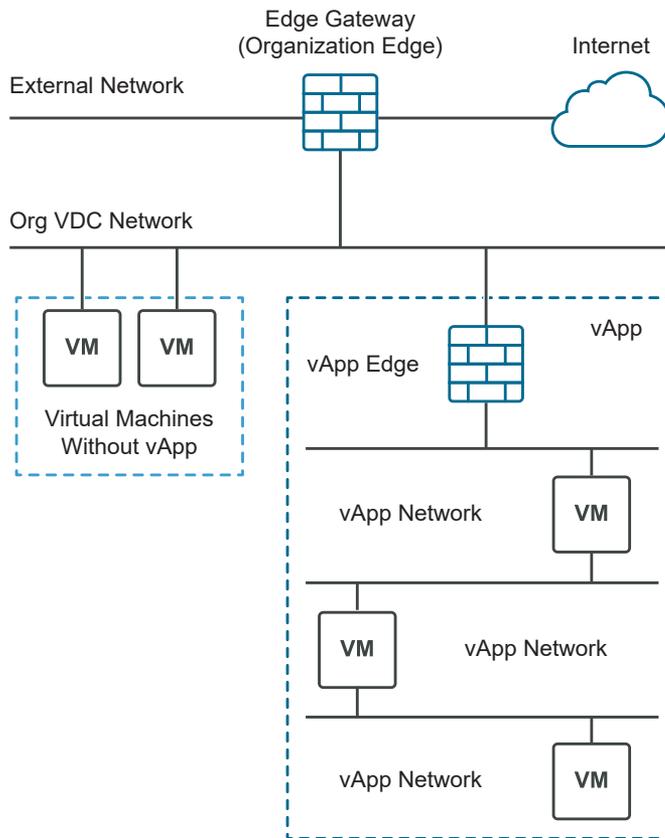
## Virtual Machine Architecture

A virtual machine can exist as a standalone machine or it can exist within a vApp.

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. The virtual machine consists of a set of specification and configuration files, and is backed by the physical resources of a host. Every virtual machine has virtual devices that provide the same functionality as physical hardware but are more portable, more secure, and easier to manage. Virtual machines can be standalone, or they can exist within a vApp. A vApp is a compound object composed of one or more virtual machines, as well as one or more networks.

The following figure shows the different options when creating a virtual machine. You can create a standalone virtual machine or a virtual machine within a vApp. The standalone virtual machine is directly connected to the organization virtual data center. You can also create a virtual machine within a vApp. By creating a virtual machine inside of a vApp, you can group together multiple virtual machines and their associated networks. vApps allow you to build complex applications, and save them to a catalog for future use.

Figure 2-1. Virtual Machines are Standalone or within a vApp



## Virtual Machine Encryption

Starting with VMware Cloud Director 10.1, you can improve the security of your data by using VM encryption. You can encrypt VMs and disks by associating them with storage policies that have the VM Encryption capability.

Encryption protects not only your virtual machine but also virtual machine disks and other files. You can view the capabilities of storage policies and the encryption status of VMs and disks in the API and UI. You can perform all operations on encrypted VMs and disks that are supported in the respective vCenter Server version.

If the organization VDC has a storage policy with enabled VM encryption, you can encrypt VMs and disks. See the [Enabling VM Encryption on Storage Policies of an Organization Virtual Data Center](#) topic in the *VMware Cloud Director Service Provider Admin Portal Guide*. To encrypt a VM or disk, associate it with a VM Encryption enabled storage policy. For virtual machines, see [Create a New Standalone Virtual Machine](#) or [Change the General Properties of a Virtual Machine](#). For named disks, see [Create a Named Disk](#) or [Edit a Named Disk](#). To decrypt a VM or disk, associate that VM or disk with a storage policy that does not have encryption enabled.

## VM Encryption Limitations

The following actions are not supported in VMware Cloud Director.

- Encrypt or decrypt a powered-on VM or its disks.
- Export an OVF of an encrypted VM.
- Encrypt and decrypt the disks of a VM with a snapshot if the disks are part of the snapshot.
- Decrypt a VM when its disk is on an encrypted policy.
- Add an encrypted disk to a non-encrypted VM.
- Encrypt an existing disk on a non-encrypted VM.
- Add an encrypted named disk to unencrypted VM.
- Create an encrypted linked clone.
- Encrypt a linked clone VM or its disks.
- Instantiate, move, or clone VMs across vCenter Server instances when the source VM is encrypted.

---

**Note** On a fast-provisioned organization VDC, if the source or target VM is encrypted and you want to create a clone, VMware Cloud Director always creates a full clone.

---

## Identifying a VM Encryption Storage Capability

By default, **System administrators** and **Organization administrators** have the necessary rights to view the organization VDC storage capabilities and whether VMs and disks are encrypted. **vApp Authors** can view the encryption status of a virtual machine and its disks on the **Details** page of the virtual machine. For more information about roles and rights, see [Predefined Roles and Their Rights](#).

## View Virtual Machines

You can view virtual machines that are standalone or part of a vApp. You can view virtual machines in a grid view or in a card view.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Choose one of the following.
  - To view the virtual machines in a grid view, click .
  - To view the virtual machines in a card view, click .

The list of virtual machines displays in a grid view or as a list of cards.

- 3 (Optional) Arrange the list of virtual machines from the **Sort by** drop-down menu.
- 4 (Optional) From the grid view, click  on the left of a virtual machine, to display the actions you can take for the selected virtual machine.  
  
For example, you can shut down a virtual machine.
- 5 To access the interface for the guest operating system of the virtual machine, click the desktop icon in the upper right corner of the card view.
- 6 To view and edit the details for a virtual machine, click **Details**.

## Create a New Standalone Virtual Machine

You can create a new standalone virtual machine.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 Click **New VM**.
- 4 Enter the name and the computer name for the virtual machine.

---

**Important** The computer name can contain only alphanumeric characters and hyphens. A computer name cannot consist of digits only and cannot contain spaces.

---

- 5 (Optional) Enter a meaningful description.
- 6 Select whether you want the virtual machine to power on right after it is created.

## 7 Select how you want to deploy the virtual machine.

Option	Action
<b>New</b>	<p>You deploy a new virtual machine with customizable settings.</p> <ul style="list-style-type: none"> <li>a Select an Operating System family and Operating System.</li> <li>b (Optional) Select a Boot image.</li> <li>c (Optional) Select a VM placement policy and a VM sizing policy.</li> </ul> <p>VM placement and VM sizing policy drop-down menus are visible only if the service provider has published such policies to the organization VDC.</p> <ul style="list-style-type: none"> <li>d (Optional) Select the size of the virtual machine from the predefined sizing options or click <b>Custom Sizing Options</b> to enter the number of virtual CPUs, cores per socket, and memory settings manually.</li> </ul> <p>If you select a VM sizing policy that defines the VM size, this option is not visible.</p> <p>The predefined sizes of the virtual machine are: <b>Small, Medium, and Large.</b></p> <ul style="list-style-type: none"> <li>e Specify the storage settings for the virtual machine, such as storage policy and size in GB.</li> <li>f Specify the network settings for the virtual machine, such as network, IP mode, IP address, and primary NIC.</li> </ul>
<b>From Template</b>	<p>You deploy a virtual machine from a template that you select from the templates catalog.</p> <ul style="list-style-type: none"> <li>a Select a virtual machine template from the list of available templates.</li> <li>b (Optional) Select a VM placement policy and a VM sizing policy.</li> </ul> <p>VM placement and VM sizing policy drop-down menus are visible only if the service provider has published such policies to the organization VDC. If the selected template has assigned policies, you might be limited to the predefined template policies.</p> <ul style="list-style-type: none"> <li>c (Optional) Select to use a custom storage policy and select the storage policy to use from the <b>Custom storage policy to use</b> drop-down menu.</li> <li>d Read and accept the end-user license agreement, if there is any.</li> </ul>

## 8 Click **OK** to save the settings of the virtual machine and to start the creation process.

You can see the card of the virtual machine in the catalog. Until the virtual machine is created, its state is displayed as **Busy**.

## Fast Provisioning of Virtual Machines

Fast provisioning saves time by using linked clones for virtual machine provisioning operations.

A linked clone is a duplicate of a virtual machine that uses the same virtual disk as the original, with a chain of delta disks to track the differences between the original and the clone. If you deactivate fast provisioning, all provisioning operations result in full clones.

A linked clone cannot exist on a different vCenter Server data center or datastore than the original virtual machine.

When you fast provision a VM, VMware Cloud Director creates a shadow virtual machine to support linked clone creation across vCenter Server data centers and datastores for the virtual machines that are associated with a specific vApp template.

A shadow virtual machine is an exact copy of the original virtual machine. The shadow virtual machine is created on the data center and datastore where the linked clone is created.

---

**Important** In-place consolidation of a fast-provisioned VM is not supported on storage containers that employ native snapshots. VVOLs and VAAI-enabled datastores use native snapshots, so fast-provisioned VMs deployed to one of these storage containers cannot be consolidated. If you need to consolidate a fast-provisioned VM deployed to a VVOL or VAAI-enabled datastore, you must relocate it to a different storage container.

---

## Opening a Virtual Machine Console

Accessing your virtual machine console allows you to view information about the virtual machine, work with the guest operating system, and perform operations that affect the guest operating system.

### Prerequisites

The virtual machine is powered on.

## Install VMware Remote Console on a Client

VMware Remote Console provides an embedded user-guest interaction in all virtual machines that are provisioned and managed by VMware Cloud Director. This section details the tasks required to install VMware Remote Console on Windows, Apple OS X, and Linux.

### Prerequisites

This operation requires the rights included in the predefined **vApp User** role or an equivalent set of rights.

### Procedure

- 1 Download the installer.
  - Navigate to the VMware Remote Console download page, and select the link for your platform.  
[www.vmware.com/go/download-vmrc](http://www.vmware.com/go/download-vmrc)
  - On the **Virtual Data Center** dashboard screen in the VMware Cloud Director Tenant Portal, click the card of the virtual data center that you want to explore. Select a virtual machine, and from the **Actions** menu select **Download VMRC**.
- 2 Run your platform installation.
  - If you are using Windows, double-click the `.msi` installer and follow the prompts.

- If you are using Linux, log in with **root** privileges, run the `.bundle` installer, and follow the prompts.
- If you are using Mac OS, double-click the `.dmg` to open it, then double-click the VMware Remote Console icon inside to copy to the Applications folder.

## Results

After installation, VMware Remote Console opens when you click Uniform Resource Identifiers (URIs) that begin with the `vmrc://` scheme. VMware Workstation, Player, and Fusion also handle the `vmrc://` URI scheme.

## Open a Virtual Machine Remote Console

You can open a virtual machine console using VMware Remote Console through the VMware Cloud Director Tenant Portal.

### Prerequisites

- Verify that VMware Remote Console is installed on your local system.
- Make sure that the selected virtual machine is in a powered-on state.
- This operation requires the rights included in the predefined **vApp User** role or an equivalent set of rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine, select **Launch VM Remote Console**.

---

**Note** If you do not have the VMware Remote Console installed, a pop-up window prompts you to either install VMware Remote Console or use the Web console.

---

## Results

The virtual machine console opens as an external virtual remote console.

---

**Note** When you connect to a VMware Cloud Director virtual machine by using VMware Remote Console, you are limited to console interaction only (sending `Ctrl+Alt+Del`). You cannot perform device operations, power operations, or settings management.

---

## Open a Web Console

You can connect to the console of a virtual machine even if you do not have VMware Remote Console installed on your local system.

### Prerequisites

- Verify that the virtual machine is powered on.
- This operation requires the rights included in the predefined **vApp User** role or an equivalent set of rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine, select **Launch Web Console**.

### Results

The virtual machine console opens in a new browser tab by using the VMware HTML Console SDK.

### What to do next

Click anywhere inside the console window to start using your mouse, keyboard, and other input devices in the console.

---

**Note** For information about supported international keyboards, see the VMware HTML Console SDK Documentation at <https://www.vmware.com/support/developer/html-console/>.

---

## Performing Power Operations on Virtual Machines

You can perform power operations on virtual machines, such as power on or off a virtual machine, suspending or resetting a virtual machine or shutting down the guest Operating System of a virtual machine.

### Power On a Virtual Machine

Powering on a virtual machine is the equivalent of powering on a physical machine.

You cannot power on a virtual machine that has guest customization enabled unless the virtual machine has a current version of VMware Tools installed.

### Prerequisites

The virtual machine is powered off.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.

- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to start, select **Power On**.

#### Results

A powered-on virtual machine displays a Powered on status in green.

## Power Off a Virtual Machine

Powering off a virtual machine is the equivalent of powering off a physical machine.

#### Prerequisites

The virtual machine is powered on.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to power off, select **Power Off**.

#### Results

A powered-off virtual machine displays a Powered off status in red.

## Shut Down a Guest Operating System

Shutting down the guest operating system of a virtual machine is the equivalent of powering off a physical machine.

#### Prerequisites

The virtual machine and guest operating system must be powered on.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine, select **Shut Down Guest OS**.

## Results

The guest OS is shut down.

## Reset a Virtual Machine

Resetting a virtual machine clears state (memory, cache, and so on), but the virtual machine continues to run. Resetting a virtual machine is the equivalent of pushing the reset button of a physical machine. It initiates a hard reset of the operating system without changing the power state of the virtual machine.

### Prerequisites

Your virtual machine is powered on.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to reset, select **Reset**.

## Results

The state clears for the virtual machine.

## Suspend a Virtual Machine

Suspending a virtual machine preserves its current state by writing the memory to disk.

The suspend and resume feature is useful when you want to save the current state of your virtual machine and continue work later from the same state.

### Prerequisites

The virtual machine is powered on.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to suspend, select **Suspend**.

## Results

The virtual machine is suspended, but its state is preserved.

## Discard the Suspended State of a Virtual Machine

If a virtual machine is in a suspended state and you no longer need to resume the use of the machine, you can discard the suspended state. Discarding the suspended state removes the saved memory and returns the machine to a powered-off state.

### Prerequisites

A virtual machine that is suspended.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine, select **Discard suspended state**.

### Results

The state is discarded, and the virtual machine is powered off.

## Power On Multiple VMs

You can power on multiple VMs simultaneously.

You cannot power on a virtual machine that has guest customization enabled unless the virtual machine has a current version of VMware Tools installed.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the VMs that you want to power on.
- 4 From the **Actions** menu, select **Power On**.
- 5 Click **OK** to confirm.

## Power Off Multiple Virtual Machines

You can power off multiple VMs simultaneously.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Toggle on the **Multiselect** option.

- 3 Select the VMs that you want to power off.
- 4 From the **Actions** menu, select **Power Off**.
- 5 Click **OK** to confirm.

## Reset Multiple Virtual Machines

Resetting multiple VMs simultaneously clears their state (memory, cache, and so on) while the VMs continue to run.

Resetting a virtual machine is the equivalent of pushing the reset button of a physical machine. It initiates a hard reset of the operating system without changing the power state of the virtual machine.

### Prerequisites

Verify that the VMs are powered on.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the VMs that you want to reset.
- 4 From the **Actions** menu, select **Reset**.
- 5 Click **OK** to confirm.

## Install VMware Tools in a Virtual Machine

VMware Cloud Director depends on VMware Tools to customize the guest OS.

VMware Tools improves management and performance of the virtual machine by replacing generic operating system drivers with VMware drivers tuned for virtual hardware. You install VMware Tools into the guest operating system. Although the guest operating system can run without VMware Tools, you lose important features and convenience.

### Prerequisites

- Verify that the virtual machine is powered on.
- If your newly created virtual machine has no guest OS, you must install it before you can install VMware Tools.
- Guest customization must be deactivated prior to installing VMware Tools.
- If the version of VMware Tools is earlier than 7299 in a virtual machine in your vApp, you must upgrade it.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine in which you want to install VMware Tools, select **Install VMware Tools**.

VMware Tools is installed on the target guest operating system. If there is an error during installation, an error message displays. You can also view the progress of the installation operation in the **Tasks** window.

- 4 To open the Web console of the virtual machine, from the **Actions** menu, select **Launch Web Console**.
- 5 Follow the instructions on the [VMware Knowledge Base Article 1014294](#) to configure the VMware Tools for your particular operating system.

## Results

VMware Tools is installed and configured on the guest operating system.

# Upgrade the Virtual Hardware Version for a Virtual Machine

You can upgrade the virtual hardware version for a virtual machine. Later virtual hardware versions support more features.

You cannot downgrade the hardware version of the virtual machines in a vApp.

VMware Cloud Director supports hardware versions depending on the backing vSphere resources. The supported hardware version depends on the latest supported virtual hardware version in the backing Provider VDC. An **Organization Administrator** or a **System Administrator** can set the hardware version to an earlier than the latest supported version by the underlying hardware. The VMware Cloud Director tenant portal dynamically sets the list of selectable virtual hardware versions based on the backing hardware of the Organization or Provider VDC.

For information about the hardware features available with virtual machine compatibility settings, see *vSphere Virtual Machine Administration*.

For information about the VMware products and their virtual hardware version, see <https://kb.vmware.com/s/article/1003746>.

## Prerequisites

- Stop the virtual machine or the vApp that contains the virtual machine.
- Verify that the latest version of VMware Tools is installed on the virtual machine.

**Procedure**

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to upgrade, select **Upgrade Virtual Hardware Version**.
- 4 Click **OK**.

**Results**

The virtual machine is upgraded to the latest version.

## Edit Virtual Machine Properties

You can edit the properties of a virtual machine, including the virtual machine name and description, hardware and network settings, guest OS settings, and so on.

### Change the General Properties of a Virtual Machine

You can review and change the name, description, and other general properties of a virtual machine.

**Prerequisites**

Changing properties such as operating system, requires that the machine is powered off.

**Procedure**

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 In the card of the virtual machine you want to edit, click **Details**.
- 4 The list of properties that you can view or edit under **General** expands by default.

Option	Action
<b>Virtual Machine Name</b>	Edit the name of the virtual machine. You can edit this property while the virtual machine is powered on.
<b>Computer Name</b>	Edit the computer and host name set in the guest operating system that identifies the virtual machine on a network. This field is restricted to 15 characters because of a Windows OS limitation on computer names. You can edit this property while the virtual machine is powered on.

Option	Action
<b>Description</b>	Edit the optional description of the virtual machine. You can edit this property while the virtual machine is powered on.
<b>Operating System Family</b>	Select an operating system family from the drop-down menu. You can edit this property while the virtual machine is powered off. In addition, you cannot edit this property if an operating system is already present on the virtual machine.
<b>Operating System</b>	Select an operating system from the drop-down menu. You can edit this property while the virtual machine is powered off. In addition, you cannot edit this property if an operating system is already present on the virtual machine.
<b>Boot Delay</b>	Specify the time in milliseconds to delay the boot operation. The time between when you power on the virtual machine and when it exits the BIOS and launches the guest operating system software can be short. You can change the boot delay to provide more time.
<b>Storage Policy</b>	Select a storage policy for the virtual machine to use from the drop-down menu. You can edit this property while the virtual machine is powered on.
<b>Virtual Data Center</b>	View the name of the virtual data center to which this virtual machine belongs.
<b>VMware Tools</b>	View whether VMware Tools is installed on the virtual machine.
<b>Virtual Hardware Version</b>	View the virtual hardware version of the virtual machine.
<b>Upgrade to:</b>	To upgrade, select a version from the drop-down menu.
<b>Synchronize time</b>	Select the check box to enable time synchronization between the virtual machine guest operating system and the virtual data center in which it is running.
<b>Enter BIOS Setup</b>	Select whether to force entry into the BIOS setup screen the next time the virtual machine boots. You can edit this property while the virtual machine is powered off.

- 5 Click **Save** once you complete making your changes.

## Change the Hardware Properties of a Virtual Machine

You can review and change the hardware properties of a virtual machine.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 In the card of the virtual machine you want to edit, click **Details**.

#### 4 Click **Hardware** to expand the list of hardware properties that you can view and edit.

Option	Description
<b>Number of virtual CPUs</b>	<p>Edit the number of CPUs.</p> <p>The maximum number of virtual CPUs that you can assign to a virtual machine depends on the number of logical CPUs on the host and the type of guest operating system that is installed on the virtual machine.</p>
<b>Cores per socket</b>	<p>Edit the cores per socket.</p> <p>You can configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single core CPU, dual-core CPU, tri-core CPU, and so on.</p>
<b>Expose hardware-assisted CPU virtualization to guest OS</b>	<p>You can expose full CPU virtualization to the guest operating system so that applications that require hardware virtualization can run on virtual machines without binary translation or paravirtualization.</p>
<b>Total Memory</b>	<p>Edit the memory resource settings for a virtual machine. The virtual machine memory size must be a multiple of 4 MB.</p> <p>This setting determines how much of the ESXi host memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size.</p>
<b>Memory hot add</b>	<p>If you enable memory hot-add, you can add memory resources to a virtual machine while the machine is powered on. This feature is only supported on certain guest operating systems and virtual machine hardware versions greater than 7.</p>
<b>Virtual CPU hot add</b>	<p>If you enable virtual CPU hot-add, you can add virtual CPUs to the virtual machine while it is powered on. You can add only multiples of the number of cores per socket. This feature is only supported on certain guest operating systems and virtual machine hardware versions.</p>
<b>Number of sockets</b>	<p>View the number of sockets.</p> <p>The number of sockets is determined by the number of virtual CPUs available. The number changes when you update the number of virtual CPUs.</p>
<b>Removable Media</b>	<p>View the available removable media, such as attached CD/DVD and floppy drive.</p>

## 5 Under **Hard Disks**, click **Add** to add a hard disk.

Option	Description
Size	Enter the hard disk size in MB. You can increase the size of the hard disk later.  <b>Note</b> You can increase the size of an existing hard disk if the virtual machine is not a linked clone and it has no snapshots.
Policy	The storage policy for the virtual machine is used by default. By default, all the hard disks attached to a virtual machine use the storage policy specified for the virtual machine. You can override this default for any of these disks when you create a virtual machine or modify its properties. The Size column for each hard disk includes a drop-down menu that lists all the storage policies available for this virtual machine.
IOPS	Select a specific IOPS for the disk. Use this option to limit the per-disk I/O operations per second.
Bus Type	Select the bus type. The options are <b>Paravirtual (SCSI)</b> , <b>LSI Logic Parallel (SCSI)</b> , <b>LSI Logic SAS (SCSI)</b> , <b>IDE</b> , and <b>SATA</b> . For more information on storage controller types and compatibility, see <i>vSphere Virtual Machine Administration Guide</i> .
Bus Number	Enter the bus number.
Unit Number	Enter the logical unit number for the hard disk drive.

## 6 Under **NICs**, click **Add** to add a new NIC.

You can add up to 10 NICs. For information about the number of supported number of NICs depending on the virtual machine hardware version, see: <http://kb.vmware.com/s/article/2051652>. VMware Cloud Director supports modifying virtual machine NICs while the virtual machine is running. For information about supported network adapter types, see <http://kb.vmware.com/kb/1001805>.

Option	Description
Primary NIC	A flag displays when the primary NIC is selected. Select a primary NIC. The primary NIC setting determines the default and only gateway for the virtual machine. The virtual machine can use any NIC to connect to virtual and physical machines that are directly connected to the same network as the NIC, but it can only use the primary NIC to connect to machines on networks that require a gateway connection.
NIC	Number of the NIC.
Connected	Select the check box to connect a NIC.
Network	Select a network from the drop-down menu.

Option	Description
IP Mode	<p>Select an IP mode.</p> <hr/> <p><b>Caution</b> Do not set the IP mode to <b>None</b> if you selected a network to which to connect the NIC.</p> <hr/> <ul style="list-style-type: none"> <li>■ <b>Static - IP Pool</b> Pulls a static IP address from the network IP pool.</li> <li>■ <b>Static - Manual</b> Allows you to specify a specific IP address manually. If you select this option, you must type an IP address in the <b>IP Address</b> column.</li> <li>■ <b>DHCP</b> Pulls an IP address from a DHCP server.</li> </ul>
MAC Address	From the drop-down menu, select whether to keep or to reset the MAC address.

7 Click **Save**.

## Change the Guest OS Customization Properties of a Virtual Machine

Guest OS customization on VMware Cloud Director is optional for all platforms. It is required for virtual machines that must join a Windows domain.

Some of the information requested on this menu applies only to Windows platforms. The Guest OS Customization panel includes the information necessary for the virtual machine to join a Windows domain. An **organization administrator** can specify default values for a domain that Windows guests in that organization can join. Not all Windows virtual machines must join a domain, but in most enterprise installations, a virtual machine that is not a domain member cannot access many of the available network resources.

### Prerequisites

- This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.
- Guest customization requires the virtual machine to be running VMware Tools.
- Before you can customize a Windows guest OS, your **system administrator** must install the appropriate Microsoft Sysprep files on the VMware Cloud Director server group. See the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.
- Customization of Linux guest operating systems requires that Perl is installed in the guest.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.

- 3 In the card of the virtual machine you want to edit, click **Details**.
- 4 Click **Guest OS Customization and Properties** to expand the list of guest operating system settings.

Option	Description
<b>Enable Guest Customization</b>	Select this option to enable guest customization.
<b>Change SID</b>	Select this option to change the Windows Security ID (SID). This option is specific for virtual machines running a Windows guest operating system. The SID is used in some Windows operating systems to uniquely identify systems and users. If you do not select this option, the new virtual machine has the same SID as the virtual machine or template on which it is based. Duplicate SIDs do not cause problems when the computers are part of a domain and only domain user accounts are used. However, if the machines are part of a Workgroup or local user accounts are used, duplicate SIDs can compromise file access controls. For more information, see the documentation for your Microsoft Windows operating system.
<b>Allow local administrator password</b>	Select this option to allow setting an administrator password on the guest operating system. <ul style="list-style-type: none"> <li>a Specify a password for the local administrator.  Leaving the <b>Specify password</b> text box blank generates a password automatically.</li> <li>b Specify the number of times to allow automatic login.  Entering a value of zero deactivates automatic login as <b>administrator</b>.</li> </ul>
<b>Require Administrators to change password on first login</b>	Select this option to require administrators to change the password of the guest operating system on the first login. This is recommended for security purposes.
<b>Auto generate password</b>	Select this option to allow password auto generation.
<b>Enable this VM to join a domain</b>	You can select this option to join the virtual machine to a Windows domain. You can use the organization's domain or override the organization's domain and enter the domain properties. <ul style="list-style-type: none"> <li>a Enter the domain name.</li> <li>b Enter the user name and password.</li> <li>c Enter the account organizational unit.</li> </ul>
<b>Script</b>	You can use a customization script to modify the guest operating system of the virtual machine. When you add a customization script to a virtual machine, the script is called only on initial customization and force recustomization. If you set the <code>precustomization</code> command line parameter, the script is called before guest customization begins. If you set the <code>postcustomization</code> command line parameter, the script is called after guest customization finishes. <ul style="list-style-type: none"> <li>■ Click the upload button below the script text box to navigate to a customization script on your local machine.</li> <li>■ Type the customization script directly into the <b>Script file</b> text box.</li> </ul> A customization script that you enter directly into the <b>Script file</b> text box cannot contain more than 1500 characters. For more information, see VMware Knowledge Base article <a href="https://kb.vmware.com/kb/1026614">https://kb.vmware.com/kb/1026614</a> .

- 5 Click **Save** once you complete making your changes.

## Understanding Guest Customization

When you customize your guest operating system, there are some settings and options you should know about.

### Enable Guest Customization Check Box

This check box is found on the **Guest OS customization** tab on the virtual machine **Properties** page. The goal of guest customization is to configure based on the options selected in the **Properties** page. If this check box is selected, guest customization and recustomization is performed when required.

This process is required for all guest customization features, such as the computer name, network settings, setting and expiring the administrator and root passwords, SID change for Windows Operating systems, and so on, to work. This option should be selected for **Power on and Force re-customization** to work.

If the check box is selected, and the virtual machine's configuration parameters in VMware Cloud Director are out of sync with the settings in the guest OS, the **Profile** tab on the virtual machines **Properties** page displays that the settings out of sync with the guest OS and the virtual machine needs guest customization.

### Guest Customization Behavior for vApps and Virtual Machines

The check boxes are deselected.

- **Enable guest customization**
- In Windows guest OSs, **Change SID**
- **Password reset**

If you want to perform a customization (or you made changes to network settings that need to be reflected in the guest OS), you can select the **Enable guest customization** check box and set the options on the **Guest OS Customization** tab of the virtual machine **Properties** page. When virtual machines from vApp templates are used to create a vApp and then add a virtual machine, the vApp templates act as building blocks. When you add virtual machines from the catalog to a new vApp, the virtual machines are enabled for guest customization by default. When you save a vApp template from a catalog as a vApp, virtual machines are enabled for guest customization only if the **Enable guest customization** check box is selected.

These are the default values of guest customization settings:

- The **Enable guest customization** check box is the same as the source virtual machine in your catalog.
- For Windows guest virtual machines, **Change SID** is the same as the source virtual machine in your catalog.
- The password reset setting is same as the source virtual machine in your catalog.

You can deselect the **Enable guest customization** check box if required before you start the vApp.

If blank virtual machines, which are pending guest OS installation, are added to a vApp, the **Enable guest customization** check box is deselected by default because these virtual machines are not yet ready for customization.

After you install the guest OS and VMware Tools, you can power off the virtual machines, stop vApp, and select the **Enable guest customization** check box and start the vApp and virtual machines to perform guest customization.

If the virtual machine name and network settings are updated on a virtual machine that has been customized, the next time you power on the virtual machine, it is recustomized, which resynchronizes the guest virtual machine with VMware Cloud Director.

## Power on and Force Recustomization of a Virtual Machine

You can power on a virtual machine and force the recustomization of a virtual machine.

If the settings in a virtual machine are not synchronized with VMware Cloud Director or an attempt to perform a guest customization has failed, you can force the recustomization of the virtual machine.

Ensure that the application that is running in the virtual machine supports a recustomization. If you change a domain controller by using Microsoft Sysprep, and also change the SID, the virtual machine might be damaged. To mitigate the risk of damaging your virtual machine, create a snapshot before you recustomize it.

### Prerequisites

- You must be an organization administrator.
- The virtual machine must be powered off.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Power** menu of the virtual machine you want to power on and customize, select **Power On and Force Recustomization**.

### Results

The virtual machine is recustomized and powered on.

## Change the Advanced Properties of a Virtual Machine

In the **Advanced** settings, you can configure the resource allocation settings (shares, reservation, and limit) to determine the amount of CPU, memory, and storage resources provided for a virtual machine.

Use the resource allocation settings (shares, reservation, and limit) to determine the amount of CPU, memory, and storage resources provided for a virtual machine.

### Resource Allocation Shares

Shares specify the relative importance of a virtual machine within a virtual data center. If a virtual machine has twice as many shares of a resource as another virtual machine, it is entitled to consume twice as much of that resource when these two virtual machines are competing for resources. Shares are typically specified as High, Normal, or Low and these values specify share values with a 4:2:1 ratio, respectively. You can also select Custom to assign a specific number of shares (which expresses a proportional weight) to each virtual machine. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

### Resource Allocation Reservation

Specifies the guaranteed minimum allocation for a virtual machine. VMware Cloud Director allows you to power on a virtual machine only if there are enough unreserved resources to satisfy the reservation of the virtual machine. The virtual data center guarantees that amount even when its resources are heavily loaded. The reservation is expressed in concrete units (megahertz or megabytes).

For example, assume that you have 2 GHz available and specify a resource allocation reservation of 1 GHz for virtual machine 1 and 1 GHz for virtual machine 2. Now each virtual machine is guaranteed to get 1 GHz if it needs it. However, if virtual machine 1 is using only 500 MHz, virtual machine 2 can use 1.5 GHz.

Reservation defaults to 0. You can specify a reservation if you need to guarantee that the minimum required amounts of CPU or memory are always available to the virtual machine.

### Resource Allocation Limit

Specifies an upper bound for CPU and memory resources that can be allocated to a virtual machine. A virtual data center can allocate more than the reservation to a virtual machine, but never allocates more than the limit, even if there are unused resources on the system. The limit is expressed in concrete units (megahertz or megabytes).

CPU and memory resource limits default to unlimited. When the memory limit is unlimited, the amount of memory configured for the virtual machine when it was created becomes its effective limit in most cases.

In most cases, it is not necessary to specify a limit. You might waste idle resources if you specify a limit. The system does not allow a virtual machine to use more resources than the limit, even when the system is underutilized, and idle resources are available. Specify a limit only if you have good reasons for doing so.

### Prerequisites

- A reservation pool virtual data center.
- Ensure that a certain amount of memory for a virtual machine is provided by the virtual data center.
- Guarantee that a particular virtual machine is always allocated a higher percentage of the virtual data center resources than other virtual machines.
- Set an upper bound on the resources that can be allocated to a virtual machine.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 In the card of the virtual machine you want to edit, click **Details**.
- 4 Click **Advanced** and **Edit**.
- 5 Set the resource allocations shares for the CPU settings by selecting an option from the **Priority** drop-down menu.

Option	Description
Low	Allocates 500 shares per virtual CPU.
Normal	Allocates 1000 shares per virtual CPU.
High	Allocates 2000 shares per virtual CPU.
Custom	Allows you to assign a specific number of shares by entering the number of shares (which expresses a proportional weight) to each virtual machine. When you assign shares to a virtual machine, you always specify the priority for that virtual machine relative to other powered-on virtual machines.

- 6 Specify the reservation for the CPU settings by entering the reservation in MHz, and optionally, the limit for the CPU settings in MHz.

Option	Description
Unlimited	The default CPU resource option.
Maximum	Specify an upper bound for CPU resources that can be allocated to a virtual machine in MHz.

- Set the resource allocations shares for the memory settings by selecting an option from the **Priority** drop-down menu.

Option	Description
Low	Allocates 5 shares per megabyte of configured virtual machine memory.
Normal	Allocates 10 shares per megabyte of configured virtual machine memory.
High	Allocates 20 shares per megabyte of configured virtual machine memory.
Custom	Allows you to assign a specific number of shares by entering the number of shares.

- Specify the reservation for the memory settings in MB and, optionally, the limit for the memory settings in MB.

Option	Description
Unlimited	The default memory resource option.
Maximum	Specify an upper bound for memory reservation that can be allocated to a virtual machine.

- Click **Save**.

## Insert Media

You can insert media such as CD/DVD images from catalogs to use in a virtual machine guest operating system. You can use these media files to install an operating system in the virtual machine, various applications, drivers, and so on.

### Prerequisites

You have access to a catalog with media files.

### Procedure

- On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- Select the virtual machine where you want to add the media.
- From the **Actions** menu, select **Insert Media**.
- On the **Insert CD** window, select the media file to insert in the virtual machine.
- Click **Insert**.

## Eject Media

After you have finished using a CD or a DVD in your virtual machine, you can eject the media file.

### Prerequisites

A media file was previously inserted to the virtual machine.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 Select the virtual machine from which you want to eject the media.
- 4 From the **Actions** menu, select **Eject Media**.

### Results

The media file is ejected.

## Copy a Virtual Machine to a Different vApp

You can copy a virtual machine to another vApp. When you copy a virtual machine, the original virtual machine remains in the source vApp.

When you copy a virtual machine, the snapshots are not included in the copy.

### Prerequisites

- This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.
- Power off the VM.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to copy, select **Copy to**.
- 4 Select the destination vApp to which you want to copy the virtual machine, and click **Next**.

- 5 Configure the resources, such as name of the virtual machine and computer name, and, optionally, the storage policy and NICs, and click **Next**.

---

**Important** The computer name can contain only alphanumeric characters and hyphens. It cannot consist of digits only and cannot contain spaces.

---

- 6 On the **Ready to Complete** page review your settings and click **Done**.

## Move a Virtual Machine to a Different vApp

You can move a virtual machine to another vApp. When you move a VM, VMware Cloud Director removes the original VM from the source vApp.

When you move a virtual machine to a different vApp, the snapshots that you have taken are lost.

Moving VMs across different vApps relies on VMware vSphere<sup>®</sup> vMotion<sup>®</sup> and Enhanced vMotion Compatibility (EVC). You can move a VM to a different vApp that belongs to the same or another organization VDC within the same organization. The organization VDC can be within the same or a different provider VDC.

While you are moving a virtual machine to a different vApp, you can perform reconfigurations such as changing the network and the storage profile.

**Table 2-1. Reconfigurations During Virtual Machine Movements and Virtual Machine States**

Reconfiguration	VM state if the target vApp is in the same organization VDC	VM state if the target vApp in another organization VDC within the same provider VDC
change the network	powered off	N/A
remove the network	powered on or off	N/A
change the storage profile	powered on or off	powered off

### Prerequisites

- Verify that you have the **vApp Author** role or an equivalent set of rights.
- Verify that the underlying vSphere resources support vMotion and EVC. For information about the requirements and limitations of vMotion and EVC, see *vCenter Server and Host Management*.
- If you want to change the VM network or the storage profile, check whether you must power off the VM. See table *Reconfigurations During VM Movements and VM States*.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.

- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the machine you want to move, select **Move to**.
- 4 Select the destination vApp, and click **Next**.
- 5 Configure the resources, such as the name of the VM and the computer name, and, optionally, the storage policy and NICs, and click **Next**.

---

**Important** The computer name can contain only alphanumeric characters and hyphens. It cannot consist of digits only and cannot contain spaces.

---

- 6 On the **Ready to Complete** page review your settings and click **Done**.

## Virtual Machine Affinity and Anti-Affinity

Affinity and anti-affinity rules allow you to spread a group of virtual machines across different ESXi hosts or keep a group of virtual machines on a particular ESXi host.

An affinity rule places a group of virtual machines on a specific host so that you can easily audit the usage of those virtual machines. An anti-affinity rule places a group of virtual machines across different hosts, which prevents all virtual machines from failing at once in the event that a single host fails.

If the affinity or anti-affinity rules cannot be satisfied, this prevents the virtual machines added to the rule from powering on.

### View Affinity and Anti-Affinity Rules

You can view existing affinity and anti-affinity rules and their properties, such as the virtual machines affected by the rules and whether the rules are enabled.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **Affinity Rules**.
- 2 (Optional) Click the **Grid editor** icon (  ) and select what details about the rules you want to be displayed.

#### Results

You see the list of the existing affinity and anti-affinity rules, virtual machines, and enabled status of each rule.

### Create an Affinity Rule

Create an affinity rule to place a specific group of virtual machines on a single host so that you can audit the usage of those virtual machines.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **Affinity Rules**.
- 2 Under **Affinity Rules**, click **New**.
- 3 Enter a name of the rule.
- 4 Deselect **Enabled** to create the rule without enabling it.  
By default, the check box is selected and rules are enabled after you create them.
- 5 Leave the **Required** check box selected.  
By default, each affinity rule is required. This means that if the rule cannot be satisfied, the virtual machines added to the rule do not power on.
- 6 Select the virtual machines that you want to add to the affinity rule.
- 7 Click **Save**.

### Results

VMware Cloud Director places the virtual machines associated with the affinity rule on a single host.

## Create an Anti-Affinity Rule

Create an anti-affinity rule to place a specific group of virtual machines across multiple hosts to prevent simultaneous failure of those virtual machines in the event that a single host fails.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **Affinity Rules**.
- 2 Under **Anti-Affinity Rules**, click **New**.
- 3 Enter a name of the rule.
- 4 Deselect **Enabled** to create the rule without enabling it.  
By default, the check box is selected and rules are enabled after you create them.
- 5 Leave the **Required** check box selected.  
By default, each anti-affinity rule is required. This means that if the rule cannot be satisfied, the virtual machines added to the rule do not power on.
- 6 Select the virtual machines to add to the anti-affinity rule.
- 7 Click **Save**.

### Results

VMware Cloud Director places the virtual machines associated with the anti-affinity rule across multiple hosts.

## Edit an Affinity or Anti-Affinity Rule

You can edit an affinity or anti-affinity rule to activate or deactivate the rule, add or remove virtual machines, change the rule name or the rule preference.

### Prerequisites

This operation requires the `Organization vDC: VM-VM Affinity Edit` right. This right is included in the predefined **Catalog Author**, **vApp Author**, and **Organization Administrator** roles.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **Affinity Rules**.
- 2 Click the radio button next to the name of the rule that you want to edit and click **Edit**.
- 3 Edit the rule properties.
  - a Change the name of the rule as necessary.
  - b Select whether to activate or deactivate the rule.
  - c Leave the **Required** check box selected.
  - d Add more virtual machines or remove virtual machines.
- 4 Click **Save**.

## Delete an Affinity or Anti-Affinity Rule

If you no longer want to use an affinity or anti affinity rule, you can delete it.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **Affinity Rules**.
- 2 Click the radio button next to the name of the rule that you want to delete and click **Delete**.
- 3 To confirm that you want to delete the rule, click **OK**.

### Results

VMware Cloud Director deletes the affinity or anti-affinity rule.

## Monitor Virtual Machines

If your VMware Cloud Director administrator has enabled the feature for monitoring virtual machines, you can view the monitoring chart from the tenant portal.

Use this feature to understand the status of a given virtual machine over time (days, weeks, or months).

## Prerequisites

This feature is only available if your VMware Cloud Director administrator has enabled it.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 Select the virtual machine you want to monitor and click **Details**.
- 4 Click **Monitoring chart** to expand the monitoring view.  
The monitoring chart displays.
- 5 Select a metric option for monitoring virtual machines.

The list in the **Metric** drop-down menu varies depending on the choices of your **system administrator**. You see some or all of the options.

Metric	Description
Disk provisioned latest	Specified in KB. Choose from day, week, or month view.
Disk read average	Specified as a percentage. Choose from day, week, or month view.
Disk write average	Specified as a percentage. Choose from day, week, or month view.
CPU usage average	Specified as a percentage. Choose from day, week, or month view.
CPU usage MHz average	Specified in MHz. Choose from day, week, or month view.
CPU usage maximum	Specified as a percentage. Choose from day, week, or month view.
Mem usage average	Specified as a percentage. Choose from day, week, or month view.
Disk used latest	Specified in KB. Choose from day, week, or month view.

A new chart is displayed each time you select a different value from the list.

- 6 (Optional) Change the time frame for metrics collection.
- 7 Click **Refresh**.
- 8 To save your changes, click **Save**.

## Working with Snapshots

Snapshots preserve the state and data of a virtual machine at the time you take the snapshot. When you take a snapshot of a virtual machine, the virtual machine is not affected and only an image of the virtual machine in a given state is copied and stored. Snapshots are useful when you must revert repeatedly to the same virtual machine state, but you do not want to create multiple virtual machines.

Snapshots are useful as a short-term solution for testing software with unknown or potentially harmful effects. For example, you can use a snapshot as a restoration point during a linear or iterative process, such as installing update packages, or during a branching process, such as installing different versions of a program.

You might want to use a snapshot when upgrading the operating system of a virtual machine. For example, before you upgrade the virtual machine, you take a snapshot to preserve the point in time before the upgrade. If there are no issues during the upgrade, you can choose to remove the snapshot, which will commit the changes you made during the upgrade. However, if you encountered an issue, you can revert to the snapshot, which will move back to your saved virtual machine state prior to the upgrade.

With VMware Cloud Director you can have only one snapshot of a virtual machine. Each attempt to take a new snapshot of a virtual machine deletes the previous one.

### Take a Snapshot of a Virtual Machine

You can take a snapshot of a virtual machine. After you take the snapshot, you can revert the virtual machine to the snapshot, or remove the snapshot.

#### Prerequisites

Verify that the virtual machine is not connected to a named disk.

---

**Note** Snapshots do not capture NIC configurations.

---

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine for which you want to take a snapshot, select **Create Snapshot**.

Taking a snapshot of a virtual machine replaces the existing snapshot if there is any.

- 4 (Optional) Select whether to snapshot the memory of the virtual machine.

When you capture the virtual machine's memory state, the snapshot retains the live state of the virtual machine. Memory snapshots create a snapshot at a precise time, for example, to upgrade software that is still working. If you take a memory snapshot and the upgrade does not complete as expected, or the software does not meet your expectations, you can revert the virtual machine to its previous state.

When you capture the memory state, the virtual machine's files do not require quiescing. If you do not capture the memory state, the snapshot does not save the live state of the virtual machine and the disks are crash consistent unless you quiesce them.

- 5 (Optional) Select whether to quiesce the guest file system.

This operation requires that VMware Tools is installed on the virtual machine. When you quiesce a virtual machine, VMware Tools quiesces the file system of the virtual machine. A quiesce operation ensures that a snapshot disk represents a consistent state of the guest file systems. Quiesced snapshots are appropriate for automated or periodic backups. For example, if you are unaware of the virtual machine's activity, but want several recent backups to revert to, you can quiesce the files.

You cannot quiesce virtual machines that have large capacity disks.

- 6 Click **OK**.

### Results

The snapshot allows you to revert your virtual machine to the most recent snapshot.

## Revert a Virtual Machine to a Snapshot

You can revert a virtual machine to the state it was in when the snapshot was created.

### Prerequisites

The virtual machine has a snapshot.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine you want to revert to a snapshot, select **Revert to Snapshot**.
- 4 Click **OK**.

### Results

The virtual machine is reverted to the saved snapshot.

## Remove a Snapshot of a Virtual Machine

You can remove a snapshot of a virtual machine.

When you remove a snapshot, you delete the state of the virtual machine that you preserved, and you can never return to that state again. Removing a snapshot does not affect the current state of the virtual machine.

### Prerequisites

A virtual machine with a stored snapshot.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine for which you want to remove the snapshot, select **Remove Snapshot**.
- 4 Click **OK**.

## Renew a Virtual Machine Lease

You can renew a virtual machine lease if the lease is expiring soon.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine with expiring lease, select **Renew Lease**.

### Results

The lease renews. You can see the new lease time frame in the **Lease** field.

## Delete a Virtual Machine

You can delete a virtual machine from your organization.

## Prerequisites

Your virtual machine must be powered off.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 From the **Actions** menu of the virtual machine that you want to delete, select **Delete**.
- 4 Confirm the deletion.

## Results

The virtual machine is deleted.

# Auto Scale Groups

Starting with VMware Cloud Director 10.2.2, you can auto scale applications depending on the current CPU and memory use.

For information about the configuration of the auto scale solution, see [Auto Scale Groups](#) in the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

Depending on predefined criteria for the CPU and memory use, VMware Cloud Director can automatically scale up or down the number of VMs in a selected scale group. To balance the load of the servers that you configure to run the same application, you can use VMware NSX Advanced Load Balancer (Avi Networks).

**System administrator** and **organization administrator** roles have full control over the VMs in the scale groups. The other global tenant roles can view the VMs and access the VM Web Console but cannot delete, edit, perform power operations, and so on.

If you delete a scale group, VMware Cloud Director does not delete any of the existing VMs in the scale group.

## Create a Scale Group

Starting with VMware Cloud Director 10.2.2, your service provider can grant you rights to create scale groups. The amount of VMs in a scale group changes automatically depending on the conditions that you define.

You can also access scale groups from a selected organization virtual data center (VDC).

## Procedure

- 1 From the top navigation bar, select **Applications** and select the **Scale Groups** tab.
- 2 Click **New Scale Group**.

- 3 Select an organization VDC in which to create the scale group.
- 4 Enter a name and, optionally, a description of the new scale group.
- 5 Select the minimum and maximum number of VMs to which you want the group to scale to, and click **Next**.
- 6 Select a VM template for the VMs in the scale group and a storage policy, and click **Next**.
- 7 Select a network for the scale group.
  - If your VDC is backed by NSX-T Data Center, select a load balancer.
  - If you want to manage the load balancer on your own or if there is no need for a load balancer, select **I have a fully set-up network**.
- 8 Click **Create Group and Add Rules**.

### Results

VMware Cloud Director starts the initial expansion of the scale group to reach the minimum number of VMs.

### What to do next

- [Add an Auto Scaling Rule](#)
- From the details view of a scale group, when you select **Monitor**, you can see all tasks related to this scale group. For example, you can see the time of creation of the scale group, all growing or shrinking tasks for the group, the rules that initiated the tasks, and so on.
- Delete a scale group. When you delete a scale group, VMware Cloud Director does not delete any of the existing VMs in the scale group. If you want to reduce the number of VMs, you must manually delete them.

## Add an Auto Scaling Rule

Starting with VMware Cloud Director 10.2.2, your service provider can grant you rights to create and manage scale groups. You can add rules that trigger the growing or shrinking of scale groups.

### Prerequisites

#### [Create a Scale Group](#)

### Procedure

- 1 From the top navigation bar, select **Applications** and select the **Scale Groups** tab.
- 2 Select a scale group and select **Rules**.
- 3 Click **Add Rule**.
- 4 Enter a name for the rule.
- 5 Select whether the scale group must expand or shrink when the rule takes effect.

- 6 Select the number of VMs by which you want the group to expand or shrink when the rule takes effect.
- 7 Enter a cooldown period in minutes after each auto scale in the group.  
The conditions cannot trigger another scaling until the cooldown period expires. The cooldown period resets when any of the rules of the scale group takes effect.
- 8 Add a condition that triggers the rule.  
The duration period is the time for which the condition must be valid to trigger the rule. To trigger the rule, all conditions must be met.
- 9 (Optional) To add another condition, click **Add Condition**.
- 10 Click **Add**.

# Working with vApps

# 3

A vApp consists of one or more virtual machines that communicate over a network and use resources and services in a deployed environment. A vApp can contain multiple virtual machines.

Starting with VMware Cloud Director 9.5, vApps support IPv6 connectivity. You can assign IPv6 addresses to virtual machines connected to IPv6 networks.

---

**Important** All steps for working with vApps are documented from the card view, assuming that you have more than one virtual data center. Completing the same procedures from the grid view is also possible, but the steps might slightly vary.

---

This chapter includes the following topics:

- [View vApps](#)
- [Build a New vApp](#)
- [Create a vApp From an OVF Package](#)
- [Add a vApp from a Catalog](#)
- [Create a vApp from a vApp Template](#)
- [Import a Virtual Machine from vCenter Server as a vApp](#)
- [Performing Power Operations on vApps](#)
- [Open a vApp](#)
- [Edit vApp Properties](#)
- [Display a vApp Network Diagram](#)
- [Working with Networks in a vApp](#)
- [Working with Snapshots](#)
- [Change the Owner of a vApp](#)
- [Move a vApp to Another Virtual Data Center](#)
- [Copy a Stopped vApp to Another Virtual Data Center](#)
- [Copy a Powered-On vApp](#)
- [Add a Virtual Machine to a vApp](#)

- [Save a vApp as a vApp Template to a Catalog](#)
- [Download a vApp as an OVF Package](#)
- [Renew a vApp Lease](#)
- [Delete a vApp](#)
- [Delete Multiple vApps](#)

## View vApps

You can view vApps in a grid view or in a card view.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 To view the vApps in a grid view, click . To view them in a card view, click . The list of vApps displays in a grid or as a list of cards.
- 3 (Optional) Configure the grid view to contain details you want to see.
  - a From the grid view, click the **Grid editor** icon (  ).
  - b Select the vApp details you want to include in the grid view by selecting the check box next to each detail you want to see.
  - c To save your changes, click **OK** .

The selected details appear as columns for each vApp.
- 4 (Optional) From the grid view, click  on the left of a vApp, to display the actions you can take for the selected vApp.

For example, you can shut down a vApp.

## Build a New vApp

Instead of creating a vApp based on a vApp template, you can decide to create a vApp using virtual machines from catalogs, new virtual machines, or a combination of both.

Building a vApp requires you to provide a name and optionally a description of the vApp. You can go back and add the virtual machines to the vApp at a later stage.

### Prerequisites

This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Select **New vApp**.
- 3 Enter a name and, optionally, a description for the vApp.
- 4 (Optional) If you want the vApp to power on upon deployment, select the **Power on** check box.

---

**Note** The vApp can power on only if there are virtual machines in it.

---

- 5 (Optional) Search the catalog for virtual machines to add to this vApp or add a new, blank virtual machine by clicking **Add Virtual Machine**.

If there are no virtual machines in the catalog, create a virtual machine and add it to the vApp.

- a Enter the name and the computer name for the virtual machine.

---

**Important** The computer name can contain only alphanumeric characters and hyphens. A computer name cannot consist of digits only and cannot contain spaces.

---

- b (Optional) Enter a meaningful description.

- c Select how you want to deploy the virtual machine.

Option	Action
<b>New</b>	<p>You deploy a new virtual machine with customizable settings.</p> <ol style="list-style-type: none"> <li>1 Select an Operating System family and Operating System.</li> <li>2 (Optional) Select a boot image.</li> <li>3 (Optional) Select a VM placement policy and a VM sizing policy.</li> </ol> <p>VM placement and VM sizing policy drop-down menus are visible only if the service provider has published such policies to the organization VDC.</p> <ol style="list-style-type: none"> <li>4 Select the size of the virtual machine or click <b>Custom Sizing Options</b> to enter the compute, memory, and storage settings manually.</li> </ol> <p>The predefined sizes of the virtual machine are small, medium, or large.</p> <ol style="list-style-type: none"> <li>5 Specify the storage options such as storage policy and size in GB.</li> <li>6 Specify the network settings for the virtual machine, such as network, IP mode, IP address, and primary NIC.</li> </ol>
<b>From Template</b>	<p>You deploy a virtual machine from a template that you select from the templates catalog.</p> <ol style="list-style-type: none"> <li>1 Select the virtual machine template from the catalog.</li> <li>2 (Optional) Select a VM placement policy and a VM sizing policy.</li> </ol> <p>VM placement and VM sizing policy drop-down menus are visible only if the service provider has published such policies to the organization VDC. If the selected template has assigned policies, you might be limited to the predefined template policies.</p> <ol style="list-style-type: none"> <li>3 (Optional) Select to use a custom storage policy and select the policy from the <b>Custom storage policy to use</b>.</li> <li>4 If there is an end-user license agreement available, you must review and accept it.</li> </ol>

- d To add the virtual machine to the vApp click **OK**.

You can see the added virtual machine in the catalog.

- 6 (Optional) Repeat [Step 5](#) for each additional virtual machine you want to create within the vApp.
- 7 To complete the creation of the vApp, click **Create**.

### Results

The vApp is created. When the vApp powers on, the virtual machines in it are created and powered on as well.

## Create a vApp From an OVF Package

You can create and deploy a vApp directly from an OVF package without creating a vApp template and a corresponding catalog item.

VMware Cloud Director has its own restrictions for OVF deployments that differ from the restrictions in vCenter Server. As a result, an OVF deployment that is successful in vCenter Server might fail in VMware Cloud Director.

VMware Cloud Director supports OVF 1.1, but it does not support all the sections of the OVF 1.1 schema. For example, the `DeploymentOptions` section in OVF is not supported.

An OVF deployment in VMware Cloud Director involves many components, such as `TransferService`, spool area on NFS mount, NFC connection to vCenter Server, checksum validation, etc. If any of these components fail, this results in OVF upload failure.

If you upload an OVF package with a manifest file, VMware Cloud Director validates the SHA-1 hash of the OVF descriptor file and all VMDK files to the values in the `manifest.mf` file. If any hash does not match, the upload fails. A **system administrator** can deactivate this check by setting the `CONFIG` property to `ovf.manifest.check.disabled`.

### Prerequisites

- Verify that you have an OVF package to upload and that you have permission to upload OVF packages and deploy vApps.
- Verify that the OVF version in the OVF descriptor file is not 0.9.
- The default maximum supported size of an OVF descriptor file in VMware Cloud Director is 12 MB. You can override this by editing the `CONFIG` property `ovf.descriptor.size.max`.
- Verify that the default maximum allowed size of the manifest file (.mf extension) is 1 MB.
- Verify that the OVF package complies with the OVF XSD schema.
- If a hardware version is provided in the `VirtualSystemType` element of the OVF descriptor file, verify that it is lower than the highest hardware version that is supported in the VDC where you upload the OVF.
- If the OVF descriptor file contains `ExtraConfig` elements, verify that your **system administrator** included these elements in `AllowedList` of `extraConfigs` elements. Elements that are not included in the `AllowedList` cause the OVF upload to fail with a validation error.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click **Add vApp from OVF**.
- 3 Click the **Upload** button and browse to a location accessible from your computer, and select the OVF/OVA template file.

The location might be your local hard drive, a network share, or a CD/DVD drive. The supported file extensions include `.ova`, `.ovf`, `.vmdk`, `.mf`, `.cert`, and `.strings` files. If you select to upload an OVF file, which references more files than you are trying to upload, for example, a VMDK file, you must browse and select all files.

- 4 Click **Next**.
- 5 Verify the details of the OVF/OVA template you are about to deploy and click **Next**.
- 6 Enter a name and, optionally a description for the vApp, and click **Next**.
- 7 (Optional) Change the computer name of the vApp so that it contains only alphanumeric characters.

This step is required only if the name of the vApp contains spaces or special characters. By default, the computer name is prepopulated with the name of the virtual machine. However, computer names must contain only alphanumeric characters.

- 8 From the **Storage Policy** drop-down menu, select a storage policy for each of the virtual machines in the vApp, and click **Next**.
- 9 Select the networks to which you want each virtual machine to connect.
  - Select a network for each virtual machine from the **Network** drop-down menu.
  - You can select the **Switch to the advanced networking workflow** check box, and enter the network settings such as primary NIC, network adapter type, network, IP assignment and IP address settings for each virtual machine in the vApp manually.

You can configure additional properties for virtual machines after you complete the wizard.

- 10 Click **Next**.
- 11 Customize the hardware of the virtual machines in the vApp, and click **Next**.

Option	Description
<b>Number of virtual CPUs</b>	Enter the number of virtual CPUs for each virtual machine in the vApp. The maximum number of virtual CPUs that you can assign to a virtual machine depends on the number of logical CPUs on the host and the type of guest operating system that is installed on the virtual machine.
<b>Cores per socket</b>	Enter the number of cores per socket for each virtual machine in the vApp. You can configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single core CPU, dual-core CPU, tri-core CPU, and so on.
<b>Number of cores</b>	View the number of cores for each virtual machine in the vApp. The number changes when you update the number of virtual CPUs.
<b>Total memory (MB)</b>	Enter the memory in MB for each virtual machine in the vApp. This setting determines how much of the ESXi host memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size.

- 12 On the Ready to Complete page, review your settings and click **Finish**.

## Results

The new vApp appears in the card view.

## Add a vApp from a Catalog

If you have access to a catalog, you can use the vApp templates in the catalog to create vApps.

A vApp template can be based on an OVF file with properties for customizing the virtual machines of the vApp. The vApp inherits these properties. If any of those properties are user-configurable, you can specify their values.

### Prerequisites

- To access vApp templates in public catalogs, verify that you are an **organization administrator** or a **vApp author**.
- To access vApp templates in organization catalogs that are shares to you, verify that you are at least a **vApp user**.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click **New** and select **Add vApp from Catalog**.
- 3 Select a template to import and click **Next**.
- 4 Enter a name and, optionally, a description for the vApp.
- 5 Enter a runtime lease and a storage lease for the vApp, and click **Next**.
- 6 From the **Storage Policy** drop-down menu, select a storage policy for each of the virtual machines in the vApp, and click **Next**.
- 7 If the placement policies and the sizing policies for the virtual machines in the vApp are configurable, select a policy for each virtual machine from the drop-down menu.
- 8 If the compute properties for the virtual machines in the vApp are configurable, customize them and click **Next**.

Option	Description
<b>Virtual CPUs</b>	Enter the number of virtual CPUs for each virtual machine in the vApp. The maximum number of virtual CPUs that you can assign to a virtual machine depends on the number of logical CPUs on the host and the type of guest operating system that is installed on the virtual machine.
<b>Cores per socket</b>	Enter the number of cores per socket for each virtual machine in the vApp. You can configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single core CPU, dual-core CPU, tri-core CPU, and so on.

Option	Description
<b>Number of cores</b>	View the number of cores for each virtual machine in the vApp. The number changes when you update the number of virtual CPUs.
<b>Memory</b>	Enter the memory in MB for each virtual machine in the vApp. This setting determines how much of the ESXi host memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size.

- 9 If the hardware properties of the virtual machines in the vApp are configurable, customize the size of the virtual machine hard disks and click **Next**.
- 10 If the networking properties of the virtual machines in the vApp are configurable, customize them and click **Next**.
  - a On the **Configure Networking** page, select the networks to which you want each virtual machine to connect.
  - b (Optional) Select the check box to switch to the advanced networking workflow and configure additional network settings for the virtual machines in the vApp.
- 11 Review the vApp settings and click **Finish**.

## Create a vApp from a vApp Template

You can create a new vApp based on a vApp template stored in a catalog to which you have access.

If the vApp template is based on an OVF file that includes OVF properties for customizing its virtual machines, those properties are passed to the vApp. If any of those properties are user-configurable, you can specify the values.

### Prerequisites

- Only organization administrators and vApp authors can access vApp templates in public catalogs.
- vApp users and above can access vApp templates in organization catalogs shared to them.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.  
The list of templates appears in a grid view.
- 2 Click the radio button next to the vApp template you want to use and click **Create vApp**.
- 3 Enter a name and, optionally, a description of the vApp.
- 4 Specify how long this vApp can run before it is automatically stopped in hours or days.

- 5 Specify for how long the stopped vApp remains available before being automatically cleaned up in hours or days.
- 6 Click **Next**.
- 7 Select the virtual data center in which you want to create the vApp.
- 8 Select a storage policy.
- 9 Click **Next**.
- 10 For VMware Cloud Director 10.2.2 and later, configure the VM placement and sizing policies.  
Starting with version 10.2.2, placement policies are global and you can publish them to multiple provider VDCs and vApp templates include both sizing and placement policy information.
- 11 Select the networks to which you want each virtual machine to connect.
  - Select a network for each virtual machine from the **Network** drop-down menu.
  - You can select the **Switch to the advanced networking workflow** check box, and enter the network settings such as primary NIC, network adapter type, network, IP assignment and IP address settings for each virtual machine in the vApp manually.

You can configure additional properties for virtual machines after you complete the wizard.
- 12 Click **Next**.
- 13 Customize the hardware of the virtual machines in the vApp, and click **Next**.

Option	Description
<b>Number of virtual CPUs</b>	Enter the number of virtual CPUs for each virtual machine in the vApp. The maximum number of virtual CPUs that you can assign to a virtual machine depends on the number of logical CPUs on the host and the type of guest operating system that is installed on the virtual machine.
<b>Cores per socket</b>	Enter the number of cores per socket for each virtual machine in the vApp. You can configure how the virtual CPUs are assigned in terms of cores and cores per socket. Determine how many CPU cores you want in the virtual machine, then select the number of cores you want in each socket, depending on whether you want a single core CPU, dual-core CPU, tri-core CPU, and so on.
<b>Number of cores</b>	View the number of cores for each virtual machine in the vApp. The number changes when you update the number of virtual CPUs.
<b>Total memory (MB)</b>	Enter the memory in MB for each virtual machine in the vApp. This setting determines how much of the ESXi host memory is allocated to the virtual machine. The virtual hardware memory size determines how much memory is available to applications that run in the virtual machine. A virtual machine cannot benefit from more memory resources than its configured virtual hardware memory size.
<b>Hard disk properties</b>	Enter the size of the virtual machine hard disk in MB.

- 14 On the Ready to Complete page, review your settings and click **Finish**.

## Results

The new vApp appears in the card view.

# Import a Virtual Machine from vCenter Server as a vApp

If you have **system administrator** rights, you can import vCenter Server VMs as vApps to VMware Cloud Director.

Importing a virtual machine does not keep the virtual machine reservation, limit, and shares settings that are configured in vCenter Server. Imported virtual machines get their resource allocation settings from the organization virtual data center on which they reside.

## Prerequisites

To see and import virtual machines from vCenter Server, verify that you have **system administrator** rights.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click **New** and select **Import from vCenter**.
- 3 From the drop-down menu, select a vCenter Server instance from which to import a virtual machine.
- 4 Select a virtual machine to import.
- 5 Enter a name and, optionally, a description for the vApp.
- 6 From the drop-down menu, select a virtual data center in which to store and run the vApp.
- 7 (Optional) From the drop-down menu, select a storage policy for the vApp.
- 8 (Optional) To delete the source virtual machine, toggle on the **Move Virtual Machine** option.
- 9 Click **Import**.

# Performing Power Operations on vApps

You can perform power operations on vApps, such as power on or off a vApp, suspending or resetting a vApp.

## Power on a vApp

Powering on a vApp powers on all the virtual machines in the vApp that are not already powered on.

## Prerequisites

You are at least a vApp author.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to power on, select **Power On**.

### Results

The vApp is powered on.

## Power off a vApp

Powering off a vApp powers off all the virtual machines in the vApp. To perform certain actions, such as adding a vApp to a catalog, copying it, or moving it to another VDC, first you must power off the vApp.

### Prerequisites

The vApp must be started.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to stop, select **Power Off**.
- 4 Click **OK**.

### Results

All virtual machines in the vApp and the vApp itself are powered off.

## Reset a vApp

Resetting a vApp clears state (memory, cache, and so on), but the vApp continues to run.

### Prerequisites

Your vApp is started and virtual machines in it are powered on.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.

- 3 From the **Actions** menu of vApp you want to reset, select **Reset**.

#### Results

The state is cleared, and the vApp continues to run.

## Suspend a vApp

Suspending a vApp preserves its current state by writing the memory to disk.

#### Prerequisites

The vApp is running.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to suspend, select **Suspend**.

#### Results

The vApp is suspended and its state is preserved.

## Discard the Suspended State of a vApp

If a vApp is in a suspended state and you no longer have to resume the use of the vApp, you can discard the suspended state. Discarding the suspended state removes the saved memory and returns the vApp to a powered-off state.

#### Prerequisites

The vApp must be in a suspended state.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the suspended vApp, select **Discard Suspended State**.

#### Results

The state is discarded and the vApp is powered off.

## Power on Multiple vApps

You can power on multiple vApps simultaneously. This action powers on all the VMs in the vApp that are not already powered on.

### Prerequisites

Verify that you are at least **vApp author**.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps that you want to power on.
- 4 From the **Actions** menu, select **Power On**.
- 5 Click **OK** to confirm.

## Power off Multiple vApps

You can power off multiple vApps simultaneously. This action powers off all the virtual machines in the vApps. To perform certain actions, such as adding a vApp to a catalog, copying it, or moving it to another virtual data center, first you must power off the vApp.

### Prerequisites

- Verify that the vApps are started.
- Verify that you are at least **vApp author**.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps that you want to power off.
- 4 From the **Actions** menu, select **Power Off**.
- 5 Click **OK** to confirm.

## Discard the Suspended State of Multiple vApps

If multiple vApps are in a suspended state and you no longer have to resume their use, you can discard the suspended state of the vApps simultaneously. Discarding the suspended state removes the saved memory and returns the vApps to a powered-off state.

### Prerequisites

- Verify that the vApps are in a suspended state.

- Verify that you are at least **vApp author**.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the suspended vApps that you want to power off.
- 4 From the **Actions** menu, select **Discard Suspended State**.

#### Results

The vApps are powered off.

## Reset Multiple vApps

Resetting multiple vApps simultaneously clears their state, which includes memory, cache, and so on, but the vApps continue to run.

#### Prerequisites

- Verify that the vApps are started and the virtual machines in them are powered on.
- Verify that you are at least **vApp author**.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps that you want to reset.
- 4 From the **Actions** menu, select **Reset** and click **OK** to confirm.

#### Results

The state of each vApp is cleared, and the vApps continue to run.

## Suspend Multiple vApps

Suspending multiple vApps simultaneously preserves their current state by writing the memory to disk.

#### Prerequisites

Verify that the vApps are running.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.

- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps that you want to suspend.
- 4 From the **Actions** menu of the vApp you want to suspend, select **Suspend**, and click **OK** to confirm.

### Results

The vApps are suspended and their state is preserved.

## Open a vApp

You can open a vApp to view the virtual machines and networks it contains. You can also view a diagram showing how the virtual machines and networks are connected.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.  
From the card view, you can see general information for each vApp, such as its name, power state, lease information, date of creation, owner, the number of virtual machines associated with the vApp, total number of CPUs, total storage and memory, and associated networks.
- 3 To view the detailed settings of a selected vApp, click **Details** on the vApp card.

## Edit vApp Properties

You can edit the properties of an existing vApp, including the vApp name and description, lease settings, order in which to start the virtual machines in the vApp, sharing settings, and network settings.

### Edit the General Properties of the vApp

You can review and change the name, description, and other general properties of a vApp.

#### Prerequisites

Verify that the vApp is powered off.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details** to view and edit the vApp properties.

- 4 Review and change the properties as needed, and click **Save**.

Option	Action
<b>Name</b>	Enter a new name for the vApp.
<b>Description</b>	Type an optional description of the vApp.
<b>Virtual data center</b>	The name of the data center to which the vApp belongs.
<b>Snapshot</b>	If there is a snapshot, details for it display.
<b>Leases</b>	<p>Select <b>Renew</b> to renew the lease.</p> <p>a Schedule the runtime lease in number of hours or days.</p> <p>Defines how long the vApp can run before it is automatically stopped.</p> <p>b Schedule the storage lease in number of hours or days.</p> <p>Defines the how long the vApp remains available before being automatically deleted.</p>

### Results

The general settings are saved.

## Edit the Start and Stop Order of Virtual Machines in a vApp

You can configure the start and stop order of virtual machines within your vApp. Configure the start and stop order in case you have applications installed in the virtual machines that must start and stop in a particular order.

These settings are useful if you need to start and stop your virtual machines in a particular order. For example, one virtual machine houses a database server, another houses an application server, and the last houses a web server. In order for the related functions to work properly, the database server must start first, the application server must start second, and the web server must start last.

### Prerequisites

Verify that the vApp is powered off.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 Click the **Start and Stop Order** tab and click **Edit**.

- 5 Edit the start and stop order properties for each virtual machine, and click **OK**.

Option	Action
<b>Start Order</b>	Enter the order in which you want the virtual machine to start. You must enter a value for each machine in the sequence.
<b>Start Action</b>	Select a start action. The start action determines what happens to a virtual machine when you start the vApp that contains it. By default, this option is set to <b>Power On</b> .
<b>Start Wait</b>	Enter the start wait time. The start wait time is the amount of time (in seconds) that you want to wait before VMware Cloud Director starts the next machine in the sequence.
<b>Stop Action</b>	Select the stop action. The stop action is the action the virtual machine takes when you stop the vApp that contains it. If you select <b>Power Off</b> , the VM powers off without performing shutdown actions that ensure stability (which is the equivalent of pulling a plug out of a socket). Select this action if you have not installed VMware Tools. Otherwise, select <b>Shut Down</b> , which ensures stability upon shutting down.
<b>Stop Wait</b>	Enter the stop wait time. The stop wait time is the amount of time (in seconds) that you want to wait before VMware Cloud Director shuts down the next virtual machine in the sequence.

## Edit the Guest Properties of a vApp

If a vApp includes user-configurable OVF properties, you can review and modify those properties.

If a virtual machine in the vApp includes a value for a user-configurable property of the same name, the virtual machine value takes precedence.

### Prerequisites

Verify that the vApp is stopped and that its guest properties are user-configurable.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and from the left panel, select **Virtual Machines**.
- 2 Click  to view the list in a card view and, optionally, arrange the list of virtual machines from the **Sort by** drop-down menu.
- 3 In the card of the virtual machine you want to edit, click **Details**.
- 4 Click **Guest Properties** and click **Edit**.
- 5 Modify the guest properties for the vApp and click **OK**.

## Share a vApp

You can share your vApps with other groups or users within your organization. The access controls that you set, determine the operations that can be completed on the shared vApps.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**, and scroll down to the sharing properties of the vApp.
- 4 Select the users with whom you want to share the vApp and click **Save**.

Option	Action
Share with everyone in the organization	<p>Select this option to share with all users in the organization and select the access level.</p> <ul style="list-style-type: none"> <li>■ To grant full control, select <b>Full Control</b>.</li> </ul> <p>All users in the organization can open, start, save a vApp as a vApp template, add the template to a catalog, change the owner of the vApp, copy to a catalog, and modify properties.</p> <ul style="list-style-type: none"> <li>■ To grant read-only access, select <b>Read Only</b>.</li> </ul>
Share with specific users and groups	<p>Select this option to share only with users that you specify.</p> <ol style="list-style-type: none"> <li>a Select the names from the <b>Users and groups with no access</b> panel to move them to the <b>Users and groups with access</b> panel.</li> <li>b Select an access level for the specified users and groups. <ul style="list-style-type: none"> <li>■ To grant full control, select <b>Full Control</b>.</li> </ul> <p>Users with full control can open, start, save a vApp as a vApp template, add the template to a catalog, change the owner of the vApp, copy to a catalog, and modify properties.</p> <ul style="list-style-type: none"> <li>■ To grant read-only access, select <b>Read Only</b>.</li> </ul> </li> </ol>

### Results

Your vApp is shared with the specified users or groups.

## Display a vApp Network Diagram

A vApp network diagram provides a graphical view of the virtual machines and networks in a vApp.

### Prerequisites

To view the vApp network diagram, your vApp must contain less than 40 virtual machines. If the vApp contains more than 40 virtual machines, the diagram is not available.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 Click the **Networks Diagram** tab.

The diagram showing how the virtual machines and the networks in the vApp are connected is displayed. A star sign represents a primary NIC. If a NIC is connected, its color is green, if a NIC is not connected, its color is white.

- 5 (Optional) To highlight the connected virtual machines and networks, click a network or a virtual machine.

The connected objects and the connections between them are highlighted.

## What to do next

You can add virtual machines or networks from this page.

## Working with Networks in a vApp

The virtual machines in a vApp can connect to vApp networks (isolated or routed) and organization virtual data center networks (direct or fenced). You can add networks of different types to a vApp to address multiple networking scenarios.

Virtual machines in the vApp can connect to the networks that are available in a vApp. If you want to connect a virtual machine to a different network, you must first add it to the vApp.

A vApp can include vApp networks and organization virtual data center networks. A vApp network can be isolated or routed. An isolated vApp network is contained within the vApp. You can also route a vApp network to an organization virtual data center network to provide connectivity to virtual machines outside of the vApp. For routed vApp networks, you can configure network services, such as a firewall and static routing.

---

**Note** Organization VDCs that are backed by NSX Data Center for vSphere support routed, isolated and direct vApp networks.

Organization VDCs that are backed by NSX-T Data Center support isolated and direct vApp networks.

---

You can connect a vApp directly to an organization virtual data center network. If you have multiple vApps that contain identical virtual machines connected to the same organization virtual data center network and you want to start the vApps at the same time, you can fence the vApp. Fencing the vApp allows you to power on the virtual machines without conflict, by isolating their MAC and IP addresses.

The networks that you add to the vApp use the network pool that is associated with the organization virtual data center in which you created the vApp.

## View vApp Networks

You can access and view the networks in a vApp.

### Prerequisites

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.

- 2 Click  to view the vApps in a card view.

- 3 In the card of the selected vApp, click **Details**.

- 4 Click the **Networks** tab.

The list of networks, if there are any, is displayed. You can view information about each network, such as name, gateway, netmask, connection and retain IP and NAT resources.

- 5 (Optional) To edit the columns to see, click the **Grid editor** icon (  ) and select or deselect the check boxes of the columns you want to be displayed or hidden, respectively.

## Fence a vApp Network

Powering on identical virtual machines which are included in different vApps might result in a conflict. To allow powering on of identical virtual machines in different vApps without conflicts, you must fence the vApp.

Fencing a vApp isolates the MAC and IP addresses of the virtual machines and changes the connection type of the organization VDC networks from direct to fenced. On the fenced networks firewall is automatically enabled and configured so that only outgoing traffic is allowed. When you fence a vApp, you can also configure NAT and firewall rules on the fenced networks.

### Prerequisites

- You can fence only direct vApp networks. If the vApp uses more than one network and the other networks are, for example, routed, only the direct network is fenced.
- The virtual machines in the vApp that use the direct network must be stopped, so that the direct vApp network is not currently in use.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.

- 2 Click  to view the vApps in a card view.

- 3 In the card of the selected vApp, click **Details**.
- 4 Click the **Networks** tab.
- 5 If the vApp is not fenced, click the **Edit** button.
- 6 Toggle on the **Fence vApp** option and click **OK**.

### Results

The IP and MAC addresses of the virtual machines become isolated. You can power on identical virtual machines in different vApps without a conflict.

## Add a Network to a vApp

You can add a network to a vApp to make the network available to the virtual machines in the vApp. You can add a vApp network or an organization virtual data center network to a vApp.

Connections can be direct or fenced. Fencing allows identical virtual machines in different vApps to be powered on without conflict by isolating the MAC and IP addresses of the virtual machines.

When fencing is enabled and the vApp is powered on, an isolated network is created from the organization virtual data center network pool. An edge gateway is created and attached to the isolated network and the organization virtual data center network. Traffic going to and from the virtual machines pass through the edge gateway, which translates the IP address using NAT and proxy-AR. This allows a router to pass traffic between two networks by using the same IP space.

### Prerequisites

To add an organization virtual data center network, your administrator must have created such a network.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Actions** and select **Add network**.

#### 4 Select the type of network to add.

Option	Action
Organization VDC Network	Select an organization virtual data center network from the list of available networks.
vApp Network	<ul style="list-style-type: none"> <li>a Enter a name and, optionally, a description for the network.</li> <li>b Enter the network gateway CIDR.</li> <li>c (Optional) Enter the primary and secondary DNS, and the DNS suffix.</li> <li>d (Optional) Select whether to allow guest VLAN.</li> <li>e (Optional) Enter static IP pool settings, such as IP ranges.</li> <li>f (Optional) To be able to connect to an organization virtual data center network, toggle on the <b>Connect to an organization VDC network</b> option and select a network from the list.</li> </ul>

#### 5 Click **Add**.

#### Results

The network is added to the vApp.

#### What to do next

Connect a virtual machine in the vApp to the network.

## Configuring Network Services for a vApp Network

You can configure network services, such as DHCP, firewalls, network address translation (NAT), and static routing for certain vApp networks.

The network services available depend on the type of vApp network.

**Table 3-1. Network Services Available by Network Type**

vApp Network Type	DHCP	Firewall	NAT	Static Routing
Direct				
Routed	X	X	X	X
Isolated	X			

**Note** Organization VDCs that are backed by NSX Data Center for vSphere support routed, isolated and direct vApp networks.

Organization VDCs that are backed by NSX-T Data Center support isolated and direct vApp networks.

## View and Edit General Network Details

You can view and edit the general vApp network details, for example the network name and description.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **General** tab, review the network information.
- 6 Click **Edit**.
- 7 Edit the vApp network name and description.
- 8 Click **Save**.

### Edit the Static IP Pool Settings of a vApp Network

You can configure a vApp network to provide static IP addresses to the virtual machines in the vApp by pulling them from a static pool of IP addresses.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **IP Management** tab, click **Static Pools**.
- 6 Click **Edit**.
- 7 Enter an IP range and click **Add**.
- 8 Click **Save**.

### Edit the DNS Settings of a vApp Network

After you create a vApp network, you can view and edit the DNS settings at any time.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.

- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **IP Management** tab, click **DNS**.  
The DNS settings are displayed.
- 6 Click **Edit**.
- 7 Edit the primary and secondary DNS, and the DNS suffix.
- 8 Click **Save**.

## Configure DHCP for a vApp Network

You can configure certain vApp networks to provide DHCP services to virtual machines in the vApp.

When you enable DHCP for a vApp network, connect a NIC on virtual machine in the vApp to that network, and select DHCP as the IP mode for that NIC. VMware Cloud Director assigns a DHCP IP address to the virtual machine when you power it on.

### Prerequisites

- Verify that the vApp network is either routed or isolated.
- Verify that the vApp is in an organization virtual data center that is backed by NSX Data Center for vSphere.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **IP Management** tab, click **DHCP**.  
The DHCP status is displayed.
- 6 Click **Edit**.
- 7 Click **Enabled**.
- 8 In the **IP Pool** text box, enter a range of IP addresses.  
VMware Cloud Director uses these addresses to satisfy DHCP requests. The range of DHCP IP addresses cannot overlap with the static IP pool for the vApp network.
- 9 Set the default and maximum lease time in seconds.
- 10 Click **Save**.

## Display the IP Allocations for Your vApp Network

You can review the IP allocations for the networks in your vApp.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **IP Management** tab, click **IP Allocations**.

The allocated IP addresses are displayed.

## Configure Static Routing for a vApp Network

You can configure certain vApp networks to provide static routing services to allow virtual machines on different vApp networks to communicate.

Any static route that you create is automatically activated.

### Prerequisites

A routed vApp network.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **Routing** tab, click **Edit**.

You can activate or deactivate static routing for the network.

## Add Static Routing for a vApp Network

You can add static routes between two vApp networks that are routed to the same organization virtual data center network. Static routes allow traffic between the networks.

You cannot add static routes to a fenced vApp or between overlapping networks. After you add a static route to a vApp network, configure the network firewall rules to allow traffic on the static route. For vApps with static routes, select to use assigned IP addresses until the vApp or associated networks are deleted.

Static routes function only when the vApps containing the routes are running. If you change the parent network of a vApp, delete a vApp, or delete a vApp network, and the vApp includes static routes, those routes cannot function and you must remove them manually.

### Prerequisites

- Two vApp networks are routed to the same organization virtual data center network.
- The vApp networks are in vApps that were started at least once.
- Static routing is enabled on both vApp networks.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 On the **Routing** tab, under Static Routing click **Add**.  
The allocated IP addresses are displayed.
- 6 Enter a name of the static route.
- 7 Enter the network address in CIDR format.  
The network address is for the vApp network to which to add a static route.
- 8 Enter the next hop IP address.  
The next hop IP address is the external IP address of that vApp network's router.
- 9 Click **Save**.
- 10 Repeat the same procedure for the second vApp network.

### Example: Static Routing Example

vApp Network 1 and vApp Network 2 are both routed to Org Network Shared. You can create a static route on each vApp network to allow traffic between the networks. You can use information about the vApp networks to create the static routes.

**Table 3-2. Network Information**

Network Name	Network Specification	Router External IP Address
vApp Network 1	192.168.1.0/24	192.168.0.100
vApp Network 2	192.168.2.0/24	192.168.0.101
Org Network Shared	192.168.0.0/24	NA

On vApp Network 1, create a static route to vApp Network 2. On vApp Network 2, create a static route to vApp Network 1.

**Table 3-3. Static Routing Settings**

vApp Network	Route Name	Network	Next Hop IP Address
vApp Network 1	tovapp2	192.168.2.0/24	192.168.0.101
vApp Network 2	tovapp1	192.168.1.0/24	192.168.0.100

## Add a Port Forwarding Rule to a vApp Network

You can configure certain vApp networks to provide port forwarding by adding a NAT mapping rule.

Port forwarding provides external access to services running on virtual machines on the vApp network.

When you configure port forwarding, VMware Cloud Director maps an external port to a service that runs on a virtual machine dedicated to inbound traffic.

When you add a port forwarding rule to a vApp network, it appears at the bottom of the NAT mapping rule list. For information about how to set the order in which port forwarding rules are enforced, see

### Prerequisites

- Verify that the vApp network is routed.
- Verify that the firewall on the vApp network is activated. If you deactivate the firewall, the NAT mapping rules are no longer applied to the vApp network.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 Click **Services** and click **Edit**.
- 6 To enable NAT, toggle on the NAT option.
- 7 From the **NAT Type** drop-down menu, select **Port Forwarding**, and click **Add**.
- 8 (Optional) To enable IP masquerading, select the check box.

- 9 Configure the port-forwarding rule.
  - a Select an external port.
  - b Select a port to which to forward.
  - c Select a virtual machine interface.
  - d Select a protocol for the type of traffic to forward.
- 10 Click **Save**.

#### What to do next

If necessary, rearrange the port-forwarding rules by using the **Move Up** or **Move Down** buttons.

## Add an IP Translation Rule to a vApp Network

You can configure certain vApp networks to provide an IP translation by adding a NAT mapping rule.

When you create an IP translation rule for a network, vCloud Director adds a DNAT and SNAT rule to the edge gateway associated with the network's port group. The DNAT rule translates an external IP address to an internal IP address for inbound traffic. The SNAT rule translates an internal IP address to an external IP address for outbound traffic.

#### Prerequisites

- Verify that the vApp network is routed.
- Verify that the firewall on the vApp network is activated. If you deactivate the firewall, the NAT mapping rules are no longer applied to the vApp network.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, click a network to view the network details.
- 5 Click **Services** and click **Edit**.
- 6 To enable NAT, toggle on the NAT option.
- 7 From the **NAT Type** drop-down menu, select **IP Translation** and click **Add**.
- 8 Select a virtual machine interface and click **Keep**.
- 9 Select a mapping mode.
- 10 If you selected **Manual** mapping mode, enter an external IP address.
- 11 Click **Save**.

### What to do next

If necessary, rearrange the IP translation rules by using the **Move Up** or **Move Down** buttons.

## Delete a vApp Network

If you no longer need a network in your vApp, you can delete the network.

### Prerequisites

The vApp is stopped and no virtual machines in the vApp are connected to the network.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 In the card of the selected vApp, click **Details**.
- 4 On the **Networks** tab, select the network that you want to delete, click **Delete**, and confirm the deletion.

## Working with Snapshots

Creating a snapshot preserves the state and data of the virtual machines within a vApp at a specific point in time. A snapshot is not intended to be used for long periods of time or instead of backing up the vApp.

You might want to use a snapshot when upgrading the virtual machines in a vApp. For example, before you upgrade the virtual machines, you create a snapshot to preserve the point in time before the upgrade. To do this, you save a snapshot prior to upgrading, and then perform the upgrade. If there are no issues during the upgrade, you can choose to remove the snapshot, which will commit the changes you made during the upgrade. However, if you encountered an issue, you can revert the snapshot, which will move back to your saved vApp state prior to the upgrade.

### Take a Snapshot of a vApp

By taking a snapshot of a vApp, you take snapshots of all virtual machines in the vApp. After you take the snapshot, you can revert all virtual machines in the vApp to the snapshot, or remove the snapshot if you do not need it.

vApp snapshots have some limitations.

- vApp snapshots do not capture NIC configurations.
- If any virtual machine in the vApp is connected to a named disk, you cannot take a vApp snapshot.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.

- 2 Click  to view the vApps in a card view.

- 3 From the **Actions** menu of the vApp for which you want to take a snapshot, select **Create Snapshot**.

Taking a snapshot of a vApp replaces the existing snapshot, if there is any.

- 4 (Optional) Select whether to snapshot the memory of the vApp.

When you capture the vApp memory state, the snapshot retains the live state of the vApp and the virtual machines in the vApp. Memory snapshots create a snapshot at a precise time, for example, to upgrade software that is still working. If you take a memory snapshot and the upgrade does not complete as expected, or the software does not meet your expectations, you can revert the virtual machine to its previous state.

When you capture the memory state, the vApp's files do not require quiescing. If you do not capture the memory state, the snapshot does not save the live state of the vApp and the disks are crash consistent unless you quiesce them.

- 5 (Optional) Select whether to quiesce the guest file system.

This operation requires that VMware Tools is installed on the virtual machines in the vApp. When you quiesce a virtual machine, VMware Tools quiesces the file system of the virtual machine. A quiesce operation ensures that a snapshot disk represents a consistent state of the guest file systems. Quiesced snapshots are appropriate for automated or periodic backups. For example, if you are unaware of the virtual machine's activity, but want several recent backups to revert to, you can quiesce the files.

You cannot quiesce vApps that have large capacity disks.

- 6 Click **OK**.

## Results

A snapshot of the vApp is created.

## What to do next

You can revert all the virtual machines in the vApp to the most recent snapshot.

## Revert a vApp to a Snapshot

You can revert all virtual machines in a vApp to the state they were in when you created the vApp snapshot.

## Prerequisites

Verify that the vApp has an existing snapshot.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to revert, select **Revert to Snapshot**.
- 4 Click **OK**.

### Results

All virtual machines in the vApp are reverted to the snapshot state.

## Remove a Snapshot of a vApp

You can remove a snapshot of a vApp.

When you remove a vApp snapshot, you delete the state of the virtual machines in the vApp snapshot and you can never return to that state again. Removing a snapshot does not affect the current state of the vApp.

### Prerequisites

You have taken a snapshot of the vApp.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp for which you want to remove a snapshot, select **Remove Snapshot**.
- 4 Click **OK**.

### Results

The snapshot is removed.

## Take Snapshots of Multiple vApps

By taking snapshots of multiple vApps, you take snapshots of all virtual machines in the vApps. After you take the snapshots, you can revert all virtual machines in the vApps to the snapshots, or remove the snapshots if you do not need them.

vApp snapshots have some limitations.

- vApp snapshots do not capture NIC configurations.

- If any virtual machine in a vApp is connected to a named disk, you cannot take a vApp snapshot.
- Taking snapshots of multiple vApps does not create snapshots of the memory of the vApps and does not quiesce the guest file system of the vApps. If you want to create a snapshot of the memory of your vApps or to quiesce the guest file system, you must create individual snapshots for each vApp. See [Take a Snapshot of a vApp](#).

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps for which you want to create snapshots.
- 4 From the **Actions** menu, select **Create Snapshot**, and click **OK** to confirm.

#### What to do next

- You can revert all the virtual machines in the vApps to the most recent snapshots. See [Revert Multiple vApps to Snapshots](#).
- You can remove the snapshots of the vApps. See [Remove the Snapshots of Multiple vApps](#).

## Remove the Snapshots of Multiple vApps

If you don't need the snapshots of multiple vApps, you can remove them simultaneously.

When you remove a vApp snapshot, you delete the state of the virtual machines in the vApp snapshot and you can never return to that state again. Removing a snapshot does not affect the current state of the vApp.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps the snapshots of which you want to remove.
- 4 From the **Actions** menu, select **Remove Snapshot**.

## Revert Multiple vApps to Snapshots

You can revert all virtual machines in multiple vApps to the state they were in when you created the vApp snapshots.

#### Prerequisites

Verify that the vApps that you want to revert have existing snapshots.

**Procedure**

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps that you want to revert to their most recent snapshots.
- 4 From the **Actions** menu, select **Revert to Snapshot**.
- 5 Click **OK** to confirm.

## Change the Owner of a vApp

You can change the owner of the vApp, for example, when a vApp owner leaves the company or changes roles within the company.

**Prerequisites**

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

**Procedure**

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp for which you want to change the owner, select **Change owner**.
- 4 Select a user from the list.
- 5 Click **OK**.

**Results**

The owner of the vApp is changed.

## Move a vApp to Another Virtual Data Center

When you move a vApp to another virtual data center, the vApp is removed from the source virtual data center.

**Prerequisites**

- You are at least a **vApp author**.
- Your vApp is powered off.

**Procedure**

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to move, select **Move to**.
- 4 Select the virtual data center where you want to move the vApp and click **OK**.
- 5 (Optional) Select the storage policy.
- 6 Click **OK**.

**Results**

The vApp is removed from the source data center and moved to the target data center.

## Copy a Stopped vApp to Another Virtual Data Center

When you copy a vApp to another virtual data center, the original vApp remains in the source virtual data center.

**Prerequisites**

- You are at least a **vApp author**.
- The vApp is powered off.

**Procedure**

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to copy, select **Copy to**.
- 4 Type a name and description.
- 5 Select the virtual data center in which you want to create the copy of the vApp.
- 6 (Optional) Select a storage policy.
- 7 Click **OK**.

**Results**

The vApp is copied with the name and description you provided to the specified virtual data center.

## Copy a Powered-On vApp

To create a vApp based on an existing vApp, you can copy a vApp and change the copy so that the copy meets your needs. You do not have to power off virtual machines in the vApp before you copy the vApp. The memory state of running virtual machines is preserved in the copied vApp.

### Prerequisites

Verify that the following conditions are met.

- You are at least a **vApp user**.
- The organization virtual data center is backed up by vCenter Server 5.5 or later.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp you want to copy, select **Copy to**.
- 4 Type a name and description.
- 5 Select the virtual data center in which you want to create the copy of the vApp.
- 6 (Optional) Select a storage policy.
- 7 Click **OK**.

### Results

A copy of the vApp is created and the vApp copy is in a suspended state. The copied vApp is enabled for network fencing.

### What to do next

Modify the network properties of the new vApp or power on the vApp.

## Add a Virtual Machine to a vApp

You can add a virtual machine to a vApp.

### Prerequisites

You must be an **organization administrator** or **vApp author** to access virtual machines in public catalogs.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.

- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp to which you want to add a virtual machine, select **Add VM**.

The list of virtual machines that are associated to the vApp displays in the **Add VMs** window.

- 4 To create a new virtual machine and associate it with the vApp automatically, click **Add Virtual Machine**.
- 5 Enter the name and the computer name for the virtual machine.

---

**Important** The computer name can contain only alphanumeric characters and hyphens. A computer name cannot consist of digits only and cannot contain spaces.

---

- 6 (Optional) Enter a meaningful description.
- 7 Select whether you want the virtual machine to power on right after it is created.
- 8 Select how you want to deploy the virtual machine.

Option	Action
<b>New</b>	<p>You deploy a new virtual machine with customizable settings.</p> <ol style="list-style-type: none"> <li>a Select an Operating System family and Operating System.</li> <li>b (Optional) Select a boot image.</li> <li>c Select the compute policy.</li> <li>d Select the size of the virtual machine or click <b>Custom Sizing Options</b> to enter the compute, memory, and storage settings manually.</li> </ol> <p>The predefined sizing options are small, medium, or large.</p> <ol style="list-style-type: none"> <li>e Specify the storage settings of the virtual machine, such as storage policy and size in GB.</li> <li>f Specify the network settings for the virtual machine, such as network, IP mode, IP address, and primary NIC.</li> </ol>
<b>From Template</b>	<p>You deploy a virtual machine from a template that you select from the templates catalog.</p> <ol style="list-style-type: none"> <li>a Select the virtual machine template from the catalog.</li> <li>b (Optional) Select to use a custom storage policy and select the policy from the <b>Custom storage policy to use</b>.</li> <li>c If there is an end user license agreement available, you must review and accept it.</li> </ol>

- 9 Click **OK** to create the virtual machine.
- 10 Click **Add** to add the virtual machine to the vApp.

## Save a vApp as a vApp Template to a Catalog

By adding a vApp to a catalog, you convert the particular vApp to a vApp template.

Starting with VMware Cloud Director 10.2.2, when you add a vApp to a catalog, the vApp template includes the placement and sizing policies of the source vApp as unmodifiable tags.

### Prerequisites

- This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.
- Your organization must have a catalog and a virtual data center with available space.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.

- 2 Click  to view the vApps in a card view.

- 3 From the **Actions** menu of the vApp you want to add to a catalog, select **Add to Catalog**.

---

**Note** You can add vApps to a catalog even if the virtual machines that belong to the vApp are in a running state. However, if you select a running vApp, it is added to the catalog as a vApp template and all the virtual machines are in a suspended state.

---

- 4 Select the destination catalog from the **Catalog** drop-down menu.
- 5 Enter a name and, optionally, a description for the vApp template.
- 6 (Optional) Select **Overwrite catalog item** if you want the new catalog item to overwrite any existing vApp template and select the catalog item to overwrite.

For example, when you upload a new version of a vApp to the catalog, you might want to overwrite the old version.

- 7 Specify how the template must be used.

The setting applies when you are creating a vApp based on the vApp template. It is ignored when you build a vApp by using individual virtual machines from this template.

Option	Description
<b>Make identical copy</b>	Select to make an identical copy of the vApp when you create a vApp from the vApp template.
<b>Customize VM settings</b>	Select to enable customization of the virtual machine settings when you create a vApp from the vApp template.

- 8 To complete the creation of the vApp template, click **OK**

### Results

The vApp template appears in the specified catalog.

## Download a vApp as an OVF Package

You can download a vApp as an OVF package or as an OVA, which is a single file distribution of the same OVF file package.

### Prerequisites

- This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.
- Verify that the vApp is powered off and undeployed.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Click  to view the vApps in a card view.
- 3 From the **Actions** menu of the vApp that you want to download, select **Download**.
- 4 Select the format in which you want to download the vApp.
- 5 (Optional) Select **Preserve identity information** to include the UUIDs and MAC addresses of the virtual machines that reside in the vApp in the downloaded OVF package.  
This limits the portability of the package and must be used only when necessary.
- 6 Click **OK** to confirm the selection and start the download.

### Results

By default, the package is downloaded in the `Downloads` folder for your browser.

## Renew a vApp Lease

If the lease of a vApp has expired, or it is about to expire, you can renew it.

### Prerequisites

Verify that you are assigned the predefined role **vApp User** or an equivalent set of rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Select the vApp you want to renew.
- 3 From the **Actions** menu, select **Renew Lease**.

- 4 Renew the runtime lease of the vApp.
  - a Select the **Runtime lease** check box.
  - b From the drop-down menu, select a value for the runtime lease.

You can select a value in hours, days, or set the lease to **Never Expires**. **System administrators** can limit the maximum length that you can choose.
- 5 Renew the storage lease of the vApp.
  - a Select the **Storage lease** check box.
  - b From the drop-down menu, select a value for the storage lease.

You can select a value in hours, days, or set the lease to **Never Expires**. **System administrators** can limit the maximum length that you can choose.

## Delete a vApp

You can delete a vApp, which removes it from your organization.

### Prerequisites

Your vApp must be stopped.

You must be at least a **vApp author**.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Select the vApp you want to delete.
- 3 From the **Actions** menu, select **Delete**.
- 4 Click **OK**.

### Results

The vApp is deleted.

## Delete Multiple vApps

To remove multiple vApps from your organization, you can delete them simultaneously.

### Prerequisites

- Verify that your vApps are stopped.
- Verify that you are at least a **vApp author**.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore, and from the left panel, select **vApps**.
- 2 Toggle on the **Multiselect** option.
- 3 Select the vApps that you want to delete.
- 4 From the **Actions** menu, select **Delete**.
- 5 To confirm, click **Delete**.

# Working with Kubernetes Clusters

# 4

You can create Kubernetes clusters of different node sizes from the existing organization VDC policies.

Kubernetes Container Clusters is the Container Service Extension plug-in for VMware Cloud Director. You can use the Kubernetes Container Clusters plug-in in the VMware Cloud Director Tenant Portal to deploy clusters with native and VMware Tanzu Kubernetes Grid Integrated Edition (TKGI) clusters. You can create Tanzu Kubernetes clusters without the Kubernetes Container Clusters plug-in.

When enabled on a vSphere cluster, VMware vSphere® with VMware Tanzu™ provides the capability to create upstream Kubernetes clusters in dedicated resource pools. For more information, see the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation.

When a service provider creates a provider VDC Kubernetes policy and publishes the policy to an organization VDC, they create an organization VDC Kubernetes policy. You can use the Kubernetes Container Clusters plug-in to create Tanzu Kubernetes clusters by applying one of the organization VDC Kubernetes policies.

## Kubernetes Runtime Options

- Tanzu Kubernetes clusters - You can use the vSphere Kubernetes runtime option to create vSphere with VMware Tanzu managed Tanzu Kubernetes clusters. This option offers more features, however, it might be more expensive. For more information, see the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation.
- Native clusters - The Kubernetes Container Clusters plug-in manages the clusters with native Kubernetes runtime. These clusters are with reduced High Availability function with a single control plane node, they offer fewer persistent volume choices and no networking automation. However, they might come at a lower cost.
- TKGI clusters - VMware Tanzu Kubernetes Grid Integrated Edition is a purpose-built container solution to operationalize Kubernetes for multi-cloud enterprises and service providers. Some of its capabilities are high availability, auto-scaling, health-checks, as well as self-healing and rolling upgrades for Kubernetes clusters. For more information on TKGI clusters, see the *VMware Tanzu Kubernetes Grid Integrated Edition* documentation.

This chapter includes the following topics:

- [Add an Organization VDC Kubernetes Policy](#)
- [Edit an Organization VDC Kubernetes Policy](#)
- [Create a Tanzu Kubernetes Cluster](#)
- [Create a Native Kubernetes Cluster](#)
- [Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster](#)
- [Configure External Access to a Service in a Tanzu Kubernetes Cluster](#)

## Add an Organization VDC Kubernetes Policy

If you have **system administrator** rights, you can add an organization VDC Kubernetes policy by using a provider VDC Kubernetes policy. You can use the organization VDC Kubernetes policy to create Tanzu Kubernetes clusters.

When you add or publish a provider VDC Kubernetes policy to an organization VDC, you make the policy available to tenants by creating an organization VDC policy. Tenants can use the available organization VDC Kubernetes policies to leverage the Kubernetes capacity while creating Tanzu Kubernetes clusters. A Kubernetes policy encapsulates placement, infrastructure quality, and persistent volume storage classes. Kubernetes policies can have different compute limits.

You can add multiple organization VDC Kubernetes policies to a single organization VDC. You can use a single provider VDC Kubernetes policy to create multiple organization VDC Kubernetes policies. You can use the organization VDC Kubernetes policies as an indicator of the service quality. For example, you can publish a Gold Kubernetes policy that allows a selection of the guaranteed machine classes and a fast storage class or a Silver Kubernetes policy that allows a selection of the best effort machine classes and a slow storage class.

### Prerequisites

- Verify that you have a **system administrator** role or a role that includes an equivalent set of rights. All other roles can only view the organization VDC Kubernetes policies.
- Verify that your environment has at least one provider VDC backed by a Supervisor Cluster. The provider VDCs backed by a Supervisor Cluster are marked with a Kubernetes icon on the **Provider VDCs** tab of the Service Provider Admin Portal. For more information on vSphere with VMware Tanzu in VMware Cloud Director, see [Using vSphere with Kubernetes in VMware Cloud Director](#) in the *VMware Cloud Director Service Provider Admin Portal Guide*.
- Verify that you are logged in to a flex organization VDC.
- Familiarize yourself with the virtual machine class types for Tanzu Kubernetes clusters. See the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation.

### Procedure

- 1 In the top navigation bar, click **Data Centers** and then click **Virtual Data Center**.

- 2 Select an organization virtual data center.
- 3 In the left panel, under **Settings**, select **Kubernetes Policies** and click **Add**.  
The **Publish to Organization VDC** wizard appears.
- 4 Enter a tenant-visible name and description for the organization VDC Kubernetes policy and click **Next**.
- 5 Select the provider VDC Kubernetes policy that you want to use and click **Next**.
- 6 Select CPU and Memory limits for the Tanzu Kubernetes clusters created under this policy.  
The maximum limits depend on the CPU and Memory allocations of the organization VDC. When you add the policy, the selected limits act as maximums for the tenants.
- 7 Choose whether you want to reserve CPU and memory for the Tanzu Kubernetes cluster nodes created in this policy and click **Next**.  
There are two editions for each class type: guaranteed and best effort. A guaranteed class edition fully reserves its configured resources, while a best effort edition allows resources to be overcommitted. Depending on your selection, on the next page of the wizard, you can select between VM class types of the guaranteed or best effort edition.
  - Select **Yes** for VM class types of the guaranteed edition for full CPU and Memory reservations.
  - Select **No** for VM class types of the best effort edition with no CPU and memory reservations.
- 8 On the **Machine classes** page of the wizard, select one or more VM class types available for this policy.  
The selected machine classes are the only class types available to tenants when you add the policy to the organization VDC.
- 9 Select one or more storage policies.
- 10 Review your choices and click **Publish**.

## Results

The information about the published policy appears in the list of Kubernetes policies. The published policy creates a Supervisor Namespace on the Supervisor Cluster with the specified resource limits from the policy.

The tenants can start using the Kubernetes policy to create Tanzu Kubernetes clusters. VMware Cloud Director places each Tanzu Kubernetes cluster created under this Kubernetes policy in the same Supervisor Namespace. The policy resource limits become resource limits for the Supervisor Namespace. All tenant-created Tanzu Kubernetes clusters in the Supervisor Namespace compete for the resources within these limits.

## What to do next

- Delete an organization VDC Kubernetes policy.

- By using the Service Provider Admin Portal, you can manage organization resource quotas. See [Manage Quotas on the Resource Consumption of an Organization](#) in the *VMware Cloud Director Service Provider Admin Portal Guide*.
- [Manage the Resource Quotas of a Group](#) or [Manage the Resource Quotas of a User](#)

## Edit an Organization VDC Kubernetes Policy

If you have **system administrator** rights, you can modify an organization VDC Kubernetes policy to change its description and the CPU and memory limits.

### Prerequisites

Verify that you have a **system administrator** role or a role that includes an equivalent set of rights. All other roles can only view the organization VDC Kubernetes policies.

### Procedure

- 1 In the top navigation bar, click **Data Centers** and then click **Virtual Data Center**.
- 2 Select an organization virtual data center.
- 3 In the left panel, under **Settings**, select **Kubernetes Policies**.
- 4 Select the organization VDC Kubernetes policy you want to edit and click **Edit**.  
The **Edit VDC Kubernetes Policy** wizard appears.
- 5 Edit the description of the organization VDC Kubernetes policy and click **Next**.  
The name of the policy is linked to the Supervisor Namespace, created during the publishing of the policy, and you cannot change it.
- 6 Edit the CPU and Memory limit for the organization VDC Kubernetes policy and click **Next**.  
You cannot edit the CPU and Memory reservation.
- 7 Review the new policy details and click **Save**.

### What to do next

- Delete an organization VDC Kubernetes policy.
- By using the Service Provider Admin Portal, you can change organization resource quotas. See [Manage Quotas on the Resource Consumption of an Organization](#) in the *VMware Cloud Director Service Provider Admin Portal Guide*.
- Change group and user quotas. See [Manage the Resource Quotas of a Group](#) or [Manage the Resource Quotas of a User](#).

## Create a Tanzu Kubernetes Cluster

You can create Tanzu Kubernetes clusters by using the Kubernetes Container Clusters plug-in.

For more information about the different Kubernetes runtime options for the cluster creation, see [Chapter 4 Working with Kubernetes Clusters](#).

You can manage Kubernetes clusters also by using the Container Service Extension CLI. See the [Container Service Extension](#) documentation.

VMware Cloud Director provisions Tanzu Kubernetes clusters with the PodSecurityPolicy Admission Controller enabled. You must create a pod security policy to deploy workloads. For information about implementing the use of pod security policies in Kubernetes, see the *Using Pod Security Policies with Tanzu Kubernetes Clusters* topic in the *vSphere with Kubernetes Configuration and Management* guide.

### Prerequisites

- Verify that your service provider published the Kubernetes Container Clusters plug-in to your organization. You can find the plug-in on the top navigation bar under **More > Kubernetes Container Clusters**.
- Verify that you have at least one organization VDC Kubernetes policy in your organization VDC. To add an organization VDC Kubernetes policy, see [Add an Organization VDC Kubernetes Policy](#).
- Verify that your service provider published the **vmware:tkgcluster Entitlement** rights bundle to your organization and granted you the **Edit: Tanzu Kubernetes Guest Cluster** right to create and modify Tanzu Kubernetes clusters. For the ability to delete clusters, you must have the **Full Control: Tanzu Kubernetes Guest Cluster** right.
- Verify that your service provider created an Access Control List (ACL) entry for you with information about your access level.

### Procedure

- 1 From the top navigation bar, select **More > Kubernetes Container Clusters**.
- 2 (Optional) If the organization VDC is enabled for TKGI cluster creation, on the **Kubernetes Container Clusters** page, select the **vSphere with Tanzu & Native** tab.
- 3 Click **New**.
- 4 Select the **vSphere with Tanzu** runtime option and click **Next**.
- 5 Enter a name for the new Kubernetes cluster and click **Next**.
- 6 Select the organization VDC to which you want to deploy a Tanzu Kubernetes cluster and click **Next**.
- 7 Select an organization VDC Kubernetes policy and a Kubernetes version, and click **Next**.

VMware Cloud Director displays a default set of Kubernetes versions that are not tied to any organization VDC or Kubernetes policy. These versions are a global setting. To change the list of available versions, use the cell management tool to run the `./cell-management-tool manage-config --name wcp.supported.kubernetes.versions -v version_numbers` command with comma-separated version numbers.

- 8 Select the number of control plane and worker nodes in the new cluster.
- 9 Select machine classes for the control plane and worker nodes, and click **Next**.
- 10 Select a Kubernetes policy storage class for the control plane and worker nodes, and click **Next**.
- 11 (Optional) For VMware Cloud Director 10.2.2 and later, specify a range of IP addresses for Kubernetes services and a range for Kubernetes pods, and click **Next**.

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

Option	Description
Pods CIDR	Specifies a range of IP addresses to use for Kubernetes pods. The default value is 192.168.0.0/16. The pods subnet size must be equal to or larger than /24. This value must not overlap with the Supervisor Cluster settings. You can enter one IP range.
Services CIDR	Specifies a range of IP addresses to use for Kubernetes services. The default value is 10.96.0.0/12. This value must not overlap with the Supervisor Cluster settings. You can enter one IP range.

- 12 Review the cluster settings and click **Finish**.

#### What to do next

- Resize the Kubernetes cluster if you want to change the number of worker nodes.
- Download the kubeconfig file. The kubectl command-line tool uses kubeconfig files to obtain information about clusters, users, namespaces, and authentication mechanisms.
- Delete a Kubernetes cluster.

## Create a Native Kubernetes Cluster

You can create Container Service Extension 3.0 managed Kubernetes clusters by using the Kubernetes Container Clusters plug-in.

For more information about the different Kubernetes runtime options for the cluster creation, see [Chapter 4 Working with Kubernetes Clusters](#).

You can manage Kubernetes clusters also by using the Container Service Extension CLI. See the [Container Service Extension](#) documentation.

#### Prerequisites

- Verify that your service provider published the Kubernetes Container Clusters plug-in to your organization. Kubernetes Container Clusters is the Container Service Extension plug-in for VMware Cloud Director. You can find the plug-in on the top navigation bar under **More > Kubernetes Container Clusters**.

- Verify that your service provider completed the Container Service Extension 3.0 server setup and published a Container Service Extension native placement policy to the organization VDC.
- Verify that your service provider published the **cse:nativeCluster Entitlement** rights bundle to your organization and granted you the **Edit CSE:NATIVECLUSTER** right to create and modify native Kubernetes clusters. For the ability to delete clusters, you must have the **Full Control CSE:NATIVECLUSTER** right.
- Verify that your service provider created an Access Control List (ACL) entry for you with information about your access level.

#### Procedure

- 1 From the top navigation bar, select **More > Kubernetes Container Clusters**.
- 2 (Optional) If the organization VDC is enabled for TKGI cluster creation, on the **Kubernetes Container Clusters** page, select the **vSphere with Tanzu & Native** tab.
- 3 Click **New**.
- 4 Select the **Native** Kubernetes runtime option.
- 5 Enter a name and select a Kubernetes Template from the list.
- 6 (Optional) Enter a description for the new Kubernetes cluster and an SSH public key.
- 7 Click **Next**.
- 8 Select the organization VDC to which you want to deploy a native cluster and click **Next**.
- 9 Select the number of control plane and worker nodes and, optionally, sizing policies for the nodes.
- 10 Click **Next**.
- 11 If you want to deploy an additional VM with NFS software, turn on the **Enable NFS** toggle.
- 12 (Optional) Select storage policies for the control plane and worker nodes.
- 13 Click **Next**.
- 14 Select a network for the Kubernetes cluster and click **Next**.
- 15 Review the cluster settings and click **Finish**.

#### What to do next

- Resize the Kubernetes cluster if you want to change the number of worker nodes.
- Download the kubeconfig file. The kubectl command-line tool uses kubeconfig files to obtain information about clusters, users, namespaces, and authentication mechanisms.
- Delete a Kubernetes cluster.

# Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster

You can create VMware Tanzu Kubernetes Grid Integrated Edition (TKGI) clusters by using the Container Service Extension.

For more information about the different Kubernetes runtime options for the cluster creation, see [Chapter 4 Working with Kubernetes Clusters](#).

You can manage Kubernetes clusters also by using the Container Service Extension CLI. See the [Container Service Extension](#) documentation.

## Prerequisites

- Verify that your service provider published the Kubernetes Container Clusters plug-in to your organization. Kubernetes Container Clusters is the Container Service Extension plug-in for VMware Cloud Director. You can find the plug-in on the top navigation bar under **More > Kubernetes Container Clusters**.
- Verify that your service provider completed the Container Service Extension 3.0 server setup and published a Container Service Extension TKGI enablement metadata to the organization VDC.
- Verify that you have the `{cse}:PKS DEPLOY RIGHT` right.

## Procedure

- 1 From the top navigation bar, select **More > Kubernetes Container Clusters**.
- 2 On the **Kubernetes Container Clusters** page, select the **TKGI** tab, and click **New**.  
The **Create New TKGI Cluster** wizard opens.
- 3 Select the organization VDC to which you want to deploy a TKGI cluster and click **Next**.  
The list might take longer to load because VMware Cloud Director requests the information from the CSE server.
- 4 Enter a name for the new TKGI cluster and select the number of worker nodes.  
TKGI clusters must have at least one worker node.
- 5 Click **Next**.
- 6 Review the cluster settings and click **Finish**.
- 7 (Optional) Click the **Refresh** button on the right side of the page for the new TKGI cluster to appear in the list of clusters.

## What to do next

- Resize the Kubernetes cluster if you want to change the number of worker nodes.
- Download the kubeconfig file. The kubectl command-line tool uses kubeconfig files to obtain information about clusters, users, namespaces, and authentication mechanisms.

- Delete a Kubernetes cluster.

## Configure External Access to a Service in a Tanzu Kubernetes Cluster

Starting with VMware Cloud Director 10.2.2, Tanzu Kubernetes clusters are by default only reachable from IP subnets of networks within the same organization virtual data center in which a cluster is created. If necessary, you can manually configure external access to specific services in a Tanzu Kubernetes cluster.

When a VDC Kubernetes policy is published to an organization VDC, a firewall policy is automatically provisioned on the cluster edge gateway to allow access to the cluster from authorized sources within the VDC. Additionally, a system SNAT rule is automatically added to the NSX-T Data Center edge gateways within the organization VDC to ensure that the cluster edge gateway is reachable by the workloads within the organization VDC.

---

**Note** If the organization virtual data center is part of an NSX-T Data Center group, the cluster edge gateway is not reachable by the other VDCs in the data center group.

---

Both the firewall policy that is provisioned on the cluster edge gateway and the SNAT rule on the NSX-T Data Center edge gateway cannot be removed unless a **system administrator** deletes the Kubernetes policy from the VDC.

If necessary, you can manually configure access from an external network to a specific service in a Tanzu Kubernetes cluster. To do that, you must create a DNAT rule on the NSX-T Data Center edge gateway which ensures that the traffic coming from external locations is forwarded to the cluster edge gateway.

### Prerequisites

- Verify that your cloud infrastructure is backed by vSphere 7.0 Update 1C, 7.0 Update 2, or later. Contact your **system administrator**.
- Verify that you are an **organization administrator**.
- Verify that your **system administrator** has created an NSX-T Data Center edge gateway within the organization virtual data center in which the Tanzu Kubernetes cluster is located.
- Verify that the public IP address that you want to use for the service was allocated to the edge gateway interface on which you want to add a DNAT rule.
- Use the `get services my-service` command of the `kubectl` command-line tool to retrieve the external IP for the service that you want to expose.

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway and, under **Services**, click **NAT**.
- 3 To add a rule, click **New**.

#### 4 Configure a DNAT rule for the service that you want to connect to an external network.

Option	Description
<b>Name</b>	Enter a meaningful name for the rule.
<b>Description</b>	(Optional) Enter a description for the rule.
<b>State</b>	To enable the rule upon creation, turn on the <b>State</b> toggle.
<b>Interface type</b>	From the drop-down menu, select DNAT.
<b>External IP</b>	Enter the public IP address of the service. The IP address that you enter must belong to the suballocated IP range of the NSX-T Data Center edge gateway.
<b>Application</b>	Leave the box empty.
<b>Internal IP</b>	Enter the service IP address that was allocated from the Kubernetes ingress pool.
<b>Internal Port</b>	(Optional) Enter a port number to which inbound traffic is directed.
<b>Logging</b>	(Optional) To have the address translation performed by this rule logged, toggle on the <b>Logging</b> option.

#### 5 Click **Save**.

#### What to do next

If you want to provide access to other applications published as Kubernetes services from external networks, you must configure additional DNAT rules for each one of them.

# Working with Networks

# 5

To provide a highly flexible and secure network infrastructure in a multipurpose cloud environment, VMware Cloud Director uses a layered networking architecture with four categories of networks. The network categories are external networks, organization virtual data center (VDC) networks, data center group networks, and vApp networks. Most types of VMware Cloud Director networks require additional infrastructure objects, such as edge gateways and network pools.

## External Networks

An external network provides an uplink interface that connects networks and virtual machines in your VMware Cloud Director environment to outside networks, such as a VPN, a corporate intranet, or the public Internet.

An external network is backed either by a single vSphere network, by multiple vSphere networks, or by an NSX-T Data Center tier-0 logical router.

Only a **system administrator** can create an external network. For information about external networks, see *VMware Cloud Director Service Provider Admin Portal Guide*.

## Network Pools

A network pool is a collection of isolated layer-2 network segments that you can use to create vApp networks and certain types of organization VDC networks on demand.

Network pools must be created before organization VDC networks and vApp networks. If they do not exist, the only network option available to an organization is the direct connection to an external network.

Only a **system administrator** can create a network pool.

For information about network pools, see *VMware Cloud Director Service Provider Admin Portal Guide*.

## Organization VDC Networks

Organization virtual data center (VDC) networks enable vApps to communicate with each other or with external networks outside the organization.

Depending on the connection of the organization VDC network to an external network, there are several different types of organization VDC networks.

Organization VDC networks provide direct or routed connections to external networks, or can be isolated from external networks and other organization VDC networks. Routed connections require an edge gateway and a network pool in the organization VDC.

A **system administrator** or an **organization administrator** creates organization VDC networks and assigns them to your organization.

A newly created organization VDC has no networks available to it. After a **system administrator** creates the required network infrastructure, an **organization administrator** can create and manage most types of organization VDC networks.

## Data Center Group Networks Backed by NSX Data Center for vSphere

A network backed by NSX Data Center for vSphere that spans a data center group. A data center group can comprise between one and 16 organization VDCs in a single or a multisite VMware Cloud Director deployment.

## Data Center Group Networks Backed by NSX-T Data Center

Data center group networks are a type of organization VDC networks that are shared between one or more VDCs and to which vApps can connect.

A **system administrator** or an **organization administrator** creates data center group networks and scopes them to a single VDC group.

VMware Cloud Director supports isolated, imported, direct, and routed data center group networks that are backed by NSX-T Data Center.

## vApp Networks

vApp networks allow virtual machines to communicate with each other or, by connecting to an organization VDC network, with virtual machines in other vApps.

A vApp network is contained within a vApp. A vApp network can be isolated from other networks or connected to an organization VDC network.

Every vApp contains a vApp network. The network is created when the vApp is deployed, and deleted when the vApp is undeployed.

An **organization administrator** sets up and controls vApp networks.

## Types of Networks in a vApp

The virtual machines in a vApp can connect to vApp networks, which can be isolated, direct, or routed, and to organization VDC networks.

---

**Note** Organization VDCs that are backed by NSX Data Center for vSphere support routed, isolated and direct vApp networks.

Organization VDCs that are backed by NSX-T Data Center support isolated and direct vApp networks.

---

You can add networks of different types to a vApp to address multiple networking scenarios.

Virtual machines in the vApp can connect to the networks that are available in a vApp. If you want to connect a virtual machine to a different network, you must first add this network to the vApp.

A vApp can include vApp networks and organization VDC networks. An isolated vApp network is contained within the vApp.

You can also route a vApp network to an organization VDC network to provide connectivity to virtual machines outside of the vApp. For routed vApp networks, you can configure network services, such as a firewall and static routing.

You can connect a vApp directly to an organization VDC network.

If you have multiple vApps that contain identical virtual machines connected to the same organization VDC network and you want to start the vApps at the same time, you can fence the vApp. Fencing the vApp allows you to power on the virtual machines without a conflict, by isolating their MAC and IP addresses.

For information, see [Working with Networks in a vApp](#).

## Edge Gateways

An edge gateway provides a routed organization VDC network with connectivity to external networks and can provide services such as load balancing, network address translation, and a firewall. VMware Cloud Director supports IPv4 and IPv6 edge gateways.

Edge gateways require NSX Data Center for vSphere or NSX-T Data Center.

This chapter includes the following topics:

- [Managing Organization Virtual Data Center Networks](#)
- [Managing Data Center Group Networking with NSX-T Data Center](#)
- [Managing Data Center Group Networking with NSX Data Center for vSphere](#)
- [Managing NSX Data Center for vSphere Edge Gateway Services](#)
- [Managing NSX-T Data Center Edge Gateways](#)

## Managing Organization Virtual Data Center Networks

A **system administrator** or an **organization administrator** creates organization VDC networks and assigns them to your organization VDC or to an organization VDC group. An **organization administrator** can view information about networks, configure network services, and more.

You can use direct, routed, isolated, or data center group organization VDC networks backed by NSX Data Center for vSphere.

You can use routed, isolated, imported, and direct organization VDC networks backed by NSX-T Data Center. You can also use routed, isolated, and imported data center group networks backed by NSX-T Data Center.

**Table 5-1. Types of Organization VDC Networks**

Data Center Type Network	Description
Direct	<p>An organization VDC network with a direct connection to one of the external networks that are provisioned by the <b>system administrator</b> and are backed by vSphere resources.</p> <p>Direct networks are supported for organization VDCs that are backed by NSX Data Center for vSphere and, starting with VMware Cloud Director 10.2.2, for organization VDCs backed by NSX-T Data Center.</p> <p>Direct networks are accessible by multiple organization VDCs.</p> <p>Virtual machines belonging to different organization VDCs can connect to and see traffic on this network.</p> <p>A direct network provides direct layer 2 connectivity to virtual machines outside of the organization VDC. Virtual machines outside of this organization VDC can connect to virtual machines in the organization VDC directly.</p> <hr/> <p><b>Note</b> Only your <b>system administrator</b> can add a direct organization VDC network.</p> <hr/> <p>Can be IPv4 or IPv6.</p>
Isolated (Internal)	<p>Isolated networks are accessible only by the same organization VDC. Only virtual machines in this organization VDC can connect to and see traffic on the internal organization VDC network.</p> <p>Isolated networks are supported for organization VDCs backed by NSX-T Data Center and for organization VDC NSX Data Center for vSphere.</p> <p>The isolated organization VDC network provides an organization VDC with an isolated, private network that multiple virtual machines and vApps can connect to. This network provides no connectivity to virtual machines outside the organization VDC. Machines outside of the organization VDC have no connectivity to machines in the organization VDC.</p>
Routed	<p>Routed networks are accessible only by the same organization VDC. Only virtual machines in this organization VDC can connect to this network.</p> <p>This network also provides controlled access to an external network. As a <b>system administrator</b> or an <b>organization administrator</b>, you can configure network address translation (NAT), firewall, and VPN settings to make specific virtual machines accessible from the external network.</p> <hr/> <p>Can be IPv4 or IPv6.</p>

Table 5-1. Types of Organization VDC Networks (continued)

Data Center Type Network	Description
Imported NSX-T Data Center Logical Switch	Imported NSX-T Data Center networks are logical segments that are created in NSX-T Data Center and use an existing NSX-T Data Center logical switch. They are imported in a specific organization as an organization VDC network.  <b>Note</b> Only a <b>system administrator</b> can import an NSX-T Data Center network.
Data Center Group Networks Backed by NSX Data Center for vSphere	This network is part of a data center group network spanning a data center group. A data center group can comprise between one and 16 organization VDCs in a single or a multisite VMware Cloud Director deployment.  Virtual machines connected to this network are connected to the underlying stretched network.
Data Center Group Networks Backed by NSX-T Data Center	Data center group networks are a type of organization VDC networks backed by NSX-T Data Center that are shared between one or more VDCs and to which vApps can connect.  Data center group networks can be isolated, imported, or routed, and require NSX-T Data Center.

All steps for managing your organization VDC networks are documented assuming that you have more than one VDC in your environment.

## View the Available Organization VDC Networks

You can view the available organization virtual data center networks.

### Prerequisites

Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.

### Procedure

- ◆ In the top navigation bar, click **Networking**.

### Results

In the **Networks** tab, you see a list of the available networks that you can filter by various criteria.

### What to do next

You can add an organization VDC network. You can also edit, increase the scope, delete, or reset an existing organization VDC network.

## Add an Isolated Organization Virtual Data Center Network

You can add an isolated organization VDC network, which is accessible only by this organization. This network provides no connectivity to virtual machines outside this organization. Virtual machines outside of this organization have no connectivity to the virtual machines in the organization.

You can add a mix of isolated and routed organization VDC networks to meet the needs of your organization. For example, you can isolate a network that contains sensitive information and have a separate network that is associated with an edge gateway and connected to the Internet.

You can create an isolated VDC network that is backed by a network pool. Your service provider can also create an isolated VDC network that is backed by an NSX-T logical switch.

You can create only an IPv4 isolated organization VDC network.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Organization Virtual Data Center**, select a VDC in which to create the network, and click **Next**.
- 4 On the **Select Network Type** page, select **Isolated** and click **Next**.
- 5 Enter a meaningful name for the network.
- 6 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.  
Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.
- 7 Enter a description of the organization VDC network.
- 8 (Optional) If the VDC in which you create the network is backed by NSX Data Center for vSphere, toggle on the **Shared** option to make the organization VDC network available to other organization VDCs within the same organization.

One potential use case for this option is when an application exists within an Organization VDC that has a reservation or allocation pool set as the allocation model. In this case, it might not have enough room to run more virtual machines. As a solution, you can create a secondary Organization VDC with pay-as-you-go and run more virtual machines on that network on a temporary basis.

---

**Note** The Organization VDCs must be backed by the same Provider VDC.

---

- 9 Click **Next**.

- 10 (Optional) To reserve one or more IP addresses for assignment to virtual machines that require static IP addresses, configure the **Static IP Pools** for the network.
  - a Enter the IP address or range of IP addresses, and click **Add**.  
To add multiple static IP addresses or ranges, repeat this step.
  - b (Optional) To modify or remove IP addresses and ranges, click **Modify** or **Remove**.
- 11 Click **Next**.
- 12 (Optional) Configure the DNS settings.

Option	Action
Primary DNS	Enter the IP address for your primary DNS server.
Secondary DNS	Enter the IP address for your secondary DNS server.
DNS Suffix	Enter your DNS suffix. The DNS suffix is the DNS name without including the host name.

- 13 Click **Next**.
- 14 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Add a Routed Organization Virtual Data Center Network

To control the access to an external network, you can add a routed organization VDC network. **System administrators** and **organization administrators** can configure network address translation (NAT), firewall, and VPN settings to make specific virtual machines accessible from the external network.

You can add a mix of routed and isolated organization VDC networks to meet the needs of your organization. For example, you can add a network that is associated with an edge gateway and connected to the Internet, while having an isolated network that contains sensitive information.

You can add an IPv4 or IPv6 routed organization VDC network.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Organization Virtual Data Center**, select a VDC in which to create the network, and click **Next**.
- 4 On the **Select Network Type** page, select **Routed** and click **Next**.
- 5 Enter a meaningful name for the network.

- 6 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.

Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.

- 7 Enter a description of the organization VDC network.

- 8 (Optional) If the VDC in which you create the network is backed by NSX Data Center for vSphere, toggle on the **Shared** option to make the organization VDC network available to other organization VDCs within the same organization.

One potential use case is when an application within an Organization VDC has a reservation or allocation pool set as the allocation model. In this case, it might not have enough room to run more virtual machines. As a solution, you can create a secondary Organization VDC with pay-as-you-go and run more virtual machines on that network on a temporary basis.

---

**Note** The Organization VDCs must share the same network pool.

---

- 9 Click **Next**.

- 10 On the **Edge Connection** page, select an edge gateway with which to associate the organization VDC network.

If the organization VDC includes more than one edge gateway, you must select an edge gateway for this network to connect to. To support another routed network, the Edge Gateway must show a value of at least 1 in the # Available Networks column.

- 11 From the **Interface Type** drop-down menu, select the interface type.

Option	Description
<b>Internal</b>	Connects to one of the Edge gateway's internal interfaces. The maximum number of networks that are allowed is 9.
<b>Distributed</b>	Creates the network on a distributed logical router connected to this edge gateway. The maximum number of networks that are allowed is 400.
<b>Subinterface</b>	Extends an organization VDC network. VMware Cloud Director identifies the network to use to extend through L2 VPN. VMware Cloud Director, with the help of NSX network virtualization, creates a trunk interface type for this network. The maximum number of networks that are allowed is 200.

- 12 (Optional) To enable tagging of guest VLANs on this network, toggle on the **Guest VLAN Allowed** option.

- 13 Click **Next**.

- 14 (Optional) To reserve one or more IP addresses for assignment to virtual machines that require static IP addresses, configure the **Static IP Pools** for the network.
  - a Enter the IP address or range of IP addresses, and click **Add**.  
To add multiple static IP addresses or ranges, repeat this step.
  - b (Optional) To modify or remove IP addresses and ranges, click **Modify** or **Remove**.
- 15 Click **Next**.
- 16 (Optional) Configure the DNS settings.

Option	Action
Primary DNS	Enter the IP address for your primary DNS server.
Secondary DNS	Enter the IP address for your secondary DNS server.
DNS Suffix	Enter your DNS suffix. The DNS suffix is the DNS name without including the host name.

- 17 Click **Next**.
- 18 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Add a Direct Organization Virtual Data Center Network

To connect to an external network by a direct route, **System administrators** can set up a direct connection.

Starting with VMware Cloud Director 10.2.2, direct network creation is supported in organization VCDs backed by NSX-T Data Center and by NSX Data Center for vSphere.

If you log in to the VMware Cloud Director Tenant Portal as an **organization administrator** and attempt to create a direct organization virtual data center network, you receive a warning message that you have insufficient rights.

### Prerequisites

Verify that you have **system administrator** rights.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Organization Virtual Data Center**, select a VDC in which to create the network, and click **Next**.
- 4 On the **Network Type** page, select **Direct** and click **Next**.
- 5 Enter a meaningful name for the network.
- 6 Enter a description of the organization VDC network.

- 7 (Optional) To make the organization VDC network available to other organization VDCs within the same organization, toggle on the **Shared** option.
- 8 On the **External Network Connection** page, select the external network to which you want your new organization virtual data center network to connect directly, and click **Next**.
- 9 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Add an Organization VDC Network with an Imported NSX-T Data Center Logical Switch

**System administrators** can create an organization VDC network by importing a logical switch from an associated NSX-T Manager instance.

### Prerequisites

- Verify that you have **system administrator** rights.
- Verify that the provider virtual data center that backs the target organization virtual data center is associated with an NSX-T Manager instance.
- You must create at least one NSX-T logical switch that is not in use by other organization virtual data center networks.

For information about creating and configuring NSX-T logical switches, see the *NSX-T Data Center Administration Guide*.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Organization Virtual Data Center**, select a VDC in which to create the network, and click **Next**.
- 4 On the **Network Type** page, select **Imported**, then select **NSX-T Logical Switch**, and click **Next**.
- 5 From the list of available NSX-T logical switches, select the target switch, and click **Next**.
- 6 Enter a meaningful name for the network.
- 7 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.  
Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.  
If the switch is configured with a subnet, this information is prepopulated.
- 8 Enter a description of the organization VDC network.
- 9 Click **Next**.

10 (Optional) Configure the DNS settings and the static IP pool.

You can add multiple IP addresses and IP ranges.

11 Click **Next**.

12 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Edit the General Settings of an Organization Virtual Data Center Network

You can modify the properties of organization VDC networks.

### Prerequisites

Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click the name of the organization VDC network that you want to edit.
- 3 On the **General** tab, click **Edit**.
  - a Edit the name and the description of the network.
  - b If the VDC in which you created the network is backed by NSX Data Center for vSphere, toggle on or off the **Shared** option to make the organization VDC network available to other organization VDCs within the same organization.
- 4 Click **Save**.

## Connect an Organization Virtual Data Center Network to an Edge Gateway

After you create an organization VDC network, you can connect the network to an edge gateway.

Starting with version 10.1, VMware Cloud Director supports connecting to an edge gateway for organization VDC networks that are backed by either NSX Data Center for vSphere or NSX-T Data Center.

### Prerequisites

This operation requires either one of the predefined **organization administrator** or **system administrator** roles, or a role that includes the **Organization VDC Network: Edit Properties** and the **VDC Group: View** right published to the organization.

### Procedure

- 1 In the top navigation bar, click **Networking**.

- 2 Click the name of the organization VDC network that you want to connect to an edge gateway.
- 3 On the **General** tab, click **Edit**.
- 4 Click **Connection**.
- 5 Connect the network to an edge gateway.
  - a Toggle on the **Connect to an edge gateway** option.
  - b Select the edge gateway to connect to from the list of available edge gateways.
  - c Select the interface type.
  - d To allow a guest VLAN, toggle the **Guest VLAN Allowed** option.
- 6 Click **Save**.

#### Results

The organization VDC network connects to an edge gateway and converts from isolated to routed.

## Disconnect an Organization VCD Network from an Edge Gateway

By disconnecting an organization VDC network from an edge gateway, you can convert it from routed to isolated.

Starting with version 10.1, connecting to and disconnecting from an edge gateway is supported for organization VDC networks that are backed by either NSX Data Center for vSphere or NSX-T Data Center.

#### Prerequisites

This operation requires either one of the predefined **organization administrator** or **system administrator** roles, or a role that includes the **Organization VDC Network: Edit Properties** right.

#### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the organization VDC network that you want to disconnect.
- 3 On the **General** tab, click **Edit**.
- 4 Click **Connection**.
- 5 To disconnect the network from the edge gateway, toggle off the **Connect to an edge gateway** option.
- 6 Click **Save**.

#### Results

You disconnected the organization VDC network from an edge gateway. The organization VDC network is converted from routed to isolated.

## Convert the Interface of a Routed Organization VDC Network

You can change the interface of a network from internal to subinterface or distributed routing, for example, by editing the network properties.

---

**Note** Cross-VDC networks cannot be converted.

---

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the organization VDC network that you want to edit.
- 3 On the **General** tab, click **Edit**.
- 4 Click **Connection**.
- 5 From the **Interface Type** drop-down menu, select the interface type.

Option	Description
<b>Internal</b>	Connects to one of the Edge gateway's internal interfaces. The maximum number of networks that are allowed is 9.
<b>Distributed</b>	Creates the network on a distributed logical router connected to this edge gateway. The maximum number of networks that are allowed is 400.
<b>Subinterface</b>	Extends an organization VDC network. VMware Cloud Director identifies the network to use to extend through L2 VPN. VMware Cloud Director, with the help of NSX network virtualization, creates a trunk interface type for this network. The maximum number of networks that are allowed is 200.

- 6 Click **Save**.

## View the IP Addresses Used for an Organization Virtual Data Center Network

You can view a list of the IP addresses from an organization virtual data center network IP pool that are currently in use.

### Prerequisites

- Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.
- Verify that your network is an isolated or routed organization virtual data center network.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network for which you want to see the used IP addresses.
- 3 In the **IP Management** section, click **IP Usage** to see which IP addresses are currently in use.

## Add IP Addresses to an Organization Virtual Data Center Network IP Pool

If an organization virtual data center network is running out of IP addresses, you can add more addresses to its IP pool.

You cannot add IP addresses to external organization virtual data center networks that have a direct connection.

### Prerequisites

- Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.
- Verify that your network is an isolated or routed organization virtual data center network.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network that you want to edit.
- 3 In the **IP Management** section, click the **Static IP Pools** tab.
- 4 Click the **Edit** button on the right.  
In the **Edit network** window, you see the gateway CIDR and the IP address ranges, if any.
- 5 In the **Static IP Pools** text box, enter the IP address or range of IP addresses and click **Add**.

---

**Note** For cross-VDC networks, the IP addresses must not overlap with the IP addresses that are assigned to the other organization VDC networks from the same stretched network.

---

- 6 Click **Save**.

### Results

The IP address or range of IP addresses are added to the network IP pool.

## Edit or Remove IP Ranges Used in an Organization Virtual Data Center Network

If an organization virtual data center network contains IP addresses that you no longer need, you can edit the addresses or delete them from the IP pool.

### Prerequisites

- Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.
- Verify that your network is an isolated or routed organization virtual data center network.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network that you want to edit.
- 3 In the **IP Management** section, click **Static IP Pools**.
- 4 Click the **Edit** button on the right.
  - To modify an IP range, select the range, make the necessary edits, and click **Modify**.
  - To remove an IP range, select the range, and click **Remove**.
- 5 Click **Save**.

## Edit the DNS Settings of an Organization Virtual Data Center Network

You can edit the DNS settings of an organization virtual data center network.

### Prerequisites

- Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.
- Verify that your network is an isolated or routed organization virtual data center network.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network that you want to edit.
- 3 In the **IP Management** section, click **DNS**.
- 4 Click the **Edit** button on the right.
- 5 Edit the primary DNS, the secondary DNS, and the DNS suffix information as necessary.
- 6 Click **Save**.

## Configure DHCP Settings for an Isolated Organization Virtual Data Center Network

You can edit the DHCP settings of an isolated organization VDC network that is backed by NSX Data Center for vSphere. The DHCP service of an organization VDC network provides IP addresses from its address pool to VM NICs that are configured to request an address from DHCP. The service provides the address when the virtual machine powers on.

Starting with version 10.2, VMware Cloud Director supports DHCP settings for both IPv4 and IPv6. You can configure IPv6 settings by using the VMware Cloud Director API.

### Prerequisites

- Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.
- Verify that your network is an isolated organization virtual data center network.
- Verify that your network is backed by NSX Data Center for vSphere.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network that you want to edit.
- 3 In the **IP Management** section, click **DHCP**.
- 4 To enable DHCP, click **Edit** on the right of **DHCP Pools Service**.
- 5 Toggle on the **DHCP Pools Service** and click **Save**.

Addresses requested by DHCP clients are pulled from a DHCP pool.

- 6 Create a DHCP pool for the network.

- a Click **New**.

- b Enter an IP address range for the pool.

The IP address range that you specify cannot overlap with the static IP address pool for the organization virtual data center.

- c Specify the default lease time for the DHCP addresses in seconds.

The default value is 3,600 seconds.

- d Specify the maximum lease time for the DHCP addresses in seconds.

This is the maximum length of time that the DHCP-assigned IP addresses are leased to the virtual machines. The default value is 7,200 seconds.

- 7 Click **Save**.

## Add a DHCP Pool to a Routed Organization Virtual Data Center Network Backed by NSX-T Data Center

You can add DHCP pools to a routed organization VDC network that is backed by NSX-T Data Center.

---

**Note** Deleting or updating DHCP pools is not supported for organization VDC networks backed by NSX-T Data Center.

---

### Prerequisites

- These operations require the predefined **organization administrator** or **system administrator** roles or a role that includes an equivalent set of rights.
- Verify that your network is a routed organization virtual data center network.
- Verify that your network is backed by NSX-T Data Center.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network that you want to edit.
- 3 In the **IP Management** section, click DHCP.
- 4 To add a DHCP pool, click **New**.
- 5 Enter an IPv4 address range for the pool.
- 6 Click **Save**.

## Edit or Delete an Existing DHCP Pool for an Isolated Organization Virtual Data Center Network Backed by NSX Data Center for vSphere

If you no longer need a DHCP pool within your isolated organization virtual data center network, you can either delete the pool that is backed by NSX Data Center for vSphere, or edit it.

### Prerequisites

- Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.
- Verify that your network is an isolated organization virtual data center network.
- Verify that the organization virtual data center network is backed by NSX Data Center for vSphere.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the name of the network that you want to edit.
- 3 Click the **IP Management** section, click **DHCP**.

#### 4 Edit or delete an existing DHCP pool.

Option	Action
Edit a DHCP pool.	<ol style="list-style-type: none"> <li>1 Select the DHCP pool that you want to edit.</li> <li>2 Click the <b>Edit</b> button.</li> <li>3 Update the IP address range for the pool.</li> <li>4 Edit the default lease time for the DHCP addresses in seconds.</li> <li>5 Edit the maximum lease time for the DHCP addresses in seconds.</li> <li>6 Click <b>Save</b>.</li> </ol>
Delete a DHCP pool.	<ol style="list-style-type: none"> <li>1 Select the DHCP pool that you want to delete.</li> <li>2 Click the <b>Delete</b> button.</li> </ol>

## Reset an Organization Virtual Data Center Network

If the network services, such as DHCP settings or firewall settings which are associated with an organization virtual data center network are not working as expected, you can reset the network.

When you reset the organization virtual data center network, you force the network DHCP service gateway to be redeployed. This operation results in a temporary disruption of the DHCP services and no network services are available while the network is resetting.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- The network is not connected to any virtual machines, vApps, or other networks.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Select an organization VDC network.
- 3 Click **Reset** and confirm the reset operation.

## Delete an Organization Virtual Data Center Network

If you no longer need an organization virtual data center network, you can delete the network.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- The network is not connected to virtual machines, vApps, or other networks.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the radio button next to the name of the target network and click **Delete**.

3 To confirm, click **OK**.

## Managing Data Center Group Networking with NSX-T Data Center

Starting with version 10.2, VMware Cloud Director supports data center group networking backed by NSX-T Data Center.

To create a network across multiple organization VDCs, you first group the VDCs and then create a group network that is shared with them.

Data center group networks backed by NSX-T Data Center provide level-2 network sharing, single active egress point configuration, and distributed firewall (DFW) rules that are applied across a data center group.

### Data center group

A data center group acts as a cross-VDC router that provides centralized networking administration, egress point configuration, and east-west traffic between all networks within the group. A data center group can contain between one and 16 VDCs that you configure to share an active egress point.

### Availability zone

An availability zone represents the compute clusters or compute fault domains that are available to the network. By default, the availability zone is the provider VDC.

---

**Important** Your **system administrator** must configure the availability zones for group networking with NSX-T Data Center by setting a **Compute Provider Scope** for the vCenter Server instance and, optionally, for the provider VDCs backed by the vCenter Server instance. By default, the compute provider scope of a provider VDC is copied from the vCenter Server instance which is backing this VDC. A **system administrator** can differentiate the compute provider scope for the different provider VDCs that are backed by a single vCenter Server instance. For example, you can have a vCenter Server instance with a scope **Germany** and a provider VDC with a scope **Munich**.

---

Your **system administrator** can also reconfigure the availability zone to be the network provider scope, which typically represents the underlying vCenter Server instance with the associated NSX-T Manager.

### Egress point

An existing NSX-T Data Center edge gateway that you configure to connect a data center group to an external network.

### Data center group network

A layer 2 network that is shared across all VDCs in a data center group.

## Managing Data Center Groups with an NSX-T Data Center Network Provider Type

After you create a data center group with an NSX-T Data Center network provider type, you can add data centers to the group, remove them, and edit the group settings.

A data center group can include up to 16 virtual data centers.

VDCs that you remove from the data center group must have no workloads attached to any of the networks that participate in the data center group.

### Create a Data Center Group with an NSX-T Data Center Network Provider Type

You can group between one and 16 VDCs in a data center group with NSX-T Data Center network provider type.

#### Prerequisites

Verify that you are an **organization administrator**, **system administrator**, or that you are assigned a role that includes an equivalent set of rights.

#### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.
- 2 Click **New**.
- 3 On the **Starting VDC** page, select a VDC backed by NSX-T Data Center to start the group.
- 4 Enter a name and, optionally, a description for the new data center group.
- 5 On the **Participating VDCs** page, select additional data centers for the new data center group, and click **Next**.
- 6 Review the data center group details and click **Finish**.

#### Results

The newly created group appears in the list of data center groups.

#### What to do next

Create a network spanning the data center group with an NSX-T Data Center network provider type.

### View and Edit the General Settings of a Data Center Group with an NSX-T Data Center Network Provider Type

You can view and edit the data center groups with an NSX-T Data Center network provider type in your organization.

#### Prerequisites

Verify that you are an **organization administrator** or that you have a role with an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 In the **General Settings** pane, click **Edit**.
- 4 Edit the name and, optionally, the description of the data center group and click **Save** to confirm.

## Manage the Participating VDCs in a Data Center Group

You can select which VDCs to be part of a VDC group and to communicate with each other.

### Prerequisites

Verify that you are an **organization administrator** or that you have a role with an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click **Participating VDCs**, and then click **Manage**.
- 4 Select the VDCs that you want to include in the group and click **Save** to confirm.

## Synchronize a Data Center Group with an NSX-T Data Center Network Provider Type

To check if all VDCs that participate in a data center group still exist and are configured properly, you can synchronize the data center group.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click **Sync** and confirm.

## Using Distributed Firewall in a Data Center Group with an NSX-T Data Center Network Provider Type

Starting with version 10.2, VMware Cloud Director supports a distributed firewall service for data center groups with an NSX-T Data Center network provider type.

When you enable a distributed firewall for a data center group with a NSX-T Data Center network provider type, you create a single default security policy that is applied to the data center group.

As an **organization administrator**, you can create and modify additional distributed firewall rules which are associated with the data center group's default security policy.

The distributed firewall service is not enabled by default. After enabling the distributed firewall, you can create IP sets and security groups to facilitate the creation of distributed firewall rules.

---

**Note** The distributed firewall rules that you create apply only to the workloads that are attached to the data center group networks.

---

## Activate Distributed Firewall for a Data Center Group with an NSX-T Data Center Network Provider Type

By using distributed firewall, you can apply a set of level 3 firewall rules across a single data center group.

Distributed firewall is not enabled by default. When you enable it, you create a single default security policy.

### Prerequisites

Verify that you are a **system administrator**.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 In the **Distributed Firewall** section, click **Activate** and confirm that you want to activate the distributed firewall.

### What to do next

Create distributed firewall rules.

## Add an IP Set to a Data Center Group

To create distributed firewall rules and add them to a data center group, you must first create IP sets. IP sets are groups of IP addresses and networks to which the distributed firewall rules apply. Combining multiple objects into IP sets helps you to reduce the total number of distributed firewall rules to be created.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.

- 2 Click the target data center group.
- 3 Under Security, click **IP Sets**.
- 4 Click **New**.
- 5 Enter a meaningful name and, optionally, a description for the new IP set.
- 6 Enter an IPv4 address, IPv6 address, or an address range in a CIDR format, and click **Add**.
- 7 To modify an existing IP address or range, click **Modify** and edit the value.
- 8 To confirm, click **Save**.

## Create a Security Group in a Data Center Group with an NSX-T Data Center Network Provider Type

Before you create distributed firewall rules for a data center group, you can group data center group networks into security groups to which the rules apply.

Security groups are groups of data center group networks to which distributed firewall rules apply. Grouping networks helps you to reduce the total number of distributed firewall rules to be created.

### Prerequisites

Verify that you have at least one data center group network that is backed by NSX-T Data Center.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Under Security, click **Security Groups** and click **New**.
- 4 Enter a name and, optionally, a description for the security group, and click **Save**.  
The new security group appears in the list.
- 5 Select the newly created security group and click **Manage Members**.
- 6 Select the data center group networks that you want to add to the security group.
- 7 Click **Save**.

### What to do next

[Add a Distributed Firewall Rule to a Data Center Group with an NSX-T Data Center Network Provider Type](#)

## Add an Application Port Profile to a Data Center Group

To create distributed firewall rules, you can use preconfigured application port profiles and custom application port profiles.

Application port profiles include a combination of a protocol and a port, or a group of ports, that is used for firewall services. In addition to the preconfigured default port profiles, you can create custom application port profiles.

#### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Under Security, click **Application Port Profiles**.
- 4 In the **Custom Applications** pane, click **New**.
- 5 Enter a name and, optionally, a description for the application port profile.
- 6 From the **Protocol** drop-down menu, select the protocol.
- 7 Enter a port, or a range of ports, separated by a comma, and click **Save**.
- 8 To configure additional port profiles, repeat the steps.

#### What to do next

Use the application port profiles to create distributed firewall rules.

## Add a Distributed Firewall Rule to a Data Center Group with an NSX-T Data Center Network Provider Type

The distributed firewall rules that you create apply only to workloads that are attached to the data center group networks.

#### Prerequisites

Verify that the distributed firewall service for the data center group is enabled.

#### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click the **Distributed Firewall** tab on the left.
- 4 Click **Edit Rules**.
- 5 To add a firewall rule, click **New on Top**.
- 6 Configure the rule.

Option	Description
Name	Enter a name for the rule.
State	To enable the rule upon creation, toggle on the <b>State</b> option.

Option	Description
<b>Applications</b>	(Optional) To select a specific port profile to which the rule applies, turn on the <b>Applications</b> toggle and click <b>Save</b> .
<b>Context</b>	(Optional) Select an NSX-T Data Center context profile for the rule.
<b>Source</b>	Select the source traffic and click <b>Keep</b> . <ul style="list-style-type: none"> <li>■ To allow or deny traffic from any source address, toggle on <b>Any Source</b>.</li> <li>■ To allow or deny traffic from specific IP sets or security groups, select the IP sets and security groups from the list.</li> </ul>
<b>Destination</b>	Select the destination traffic and click <b>Keep</b> . <ul style="list-style-type: none"> <li>■ To allow or deny traffic to any destination address, toggle on <b>Any Destination</b>.</li> <li>■ To allow or deny traffic to specific IP sets or security groups, select the IP sets and security groups from the list.</li> </ul>
<b>Action</b>	From the <b>Action</b> drop-down menu, select whether to allow or deny traffic from or to specific sources. <ul style="list-style-type: none"> <li>■ To allow traffic from or to the specified sources, destinations, and services, select <b>Accept</b>.</li> <li>■ To block traffic from or to the specified sources, destinations, and services, select <b>Deny</b>.</li> </ul>
<b>IP Protocol</b>	Select whether to apply the rule to IPv4 or IPv6 traffic.
<b>Enable logging.</b>	To have the address translation performed by this rule logged, turn on the <b>Enable logging</b> toggle.

7 Click **Save**.

8 To configure additional rules, repeat the steps.

## Results

After you create the firewall rules, they appear in the Distributed Firewall Rules list. You can move the rules up or down, edit, or delete the rules, as needed.

## Deactivate the Default Distributed Firewall Policy

If you want to deactivate the distributed firewall service, you must first deactivate the default distributed firewall policy.

When you deactivate the default policy, you can edit the distributed firewall rules, but the rules are no longer applied.

## Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click the **Distributed Firewall** tab on the left.

- 4 In the **Default Policy** card above the distributed firewall rules list, click **Disable** and confirm the action.

### Results

The default policy is deactivated. The rest of the distributed firewall rules can be edited but they are not applied.

## Deactivate the Distributed Firewall Service

If you do not want to use the distributed firewall service, you can deactivate it.

When you deactivate the distributed firewall service for a data center group, the security rules configuration for this group is deleted permanently and cannot be recovered.

### Prerequisites

[Deactivate the Default Distributed Firewall Policy](#)

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click **General**.
- 4 In the **Distributed Firewall** pane on the right, click **Deactivate**, and confirm the action.

### Results

The distributed firewall service is deactivated and the security rules configuration is deleted.

## Managing Data Center Group Networks with an NSX-T Data Center Network Provider Type

After you create and configure a data center group, you can create and manage data center group networks spanning the participating VDCs.

You can use routed, isolated, and imported organization data center group networks backed by NSX-T Data Center.

A data center group network can only be scoped to a single data center group.

You can increase the scope of an existing network from an organization VDC to a data center group.

You can add all types of networks to a data center group.

---

**Important** The IP addresses in the networks that participate in a data center group must not overlap, even if the networks are isolated.

---

Table 5-2. Types of Data Center Group Networks

Data Center Group Network Type	Description
Isolated	An isolated data center group network is accessible only by VDCs in the same data center group. Only virtual machines in the data center group can connect to and see traffic on the isolated data center group network.
Routed	A routed data center group network provides controlled access to an external network through an NSX-T Data Center edge gateway that is part of the data center group.
Imported	An imported data center group network uses an existing NSX-T Data Center logical switch. Only a <b>system administrator</b> can import a network.

## Create an Isolated Data Center Group Network Backed by an NSX-T Data Center

You can add an isolated data center group network, which is accessible only to VMs in the data center group. VMs outside of this network have no connectivity to it, regardless of whether they are connected to other networks in the same data center group.

### Prerequisites

- Verify that you are an **organization administrator**.
- Verify that you have created a data center group with an NSX-T Data Center network provider type.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Data Center Group**, and select a group with an NSX-T Data Center network provider in which to create the network.
- 4 On the **Network Type** page, select **Isolated** and click **Next**.
- 5 Enter a meaningful name for the network.
- 6 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.  
Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.
- 7 Enter a description of the organization VDC network.
- 8 Click **Next**.
- 9 (Optional) To reserve one or more IP addresses for assignment to virtual machines that require static IP addresses, configure the **Static IP Pools** for the network.
  - a Enter the IP address or range of IP addresses, and click **Add**.  
To add multiple static IP addresses or ranges, repeat this step.
  - b (Optional) To modify or remove IP addresses and ranges, click **Modify** or **Remove**.

## 10 (Optional) Configure the DNS settings.

Option	Action
Primary DNS	Enter the IP address for your primary DNS server.
Secondary DNS	Enter the IP address for your secondary DNS server.
DNS Suffix	Enter your DNS suffix. The DNS suffix is the DNS name without including the host name.

11 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Create a Routed Data Center Group Network Backed by NSX-T Data Center

To control the access to an external network, you can add a routed data center group network.

### Prerequisites

- Verify that you are an **organization administrator** or that you have a role with an equivalent set of rights.
- Verify that you have created a data center group with an NSX-T Data Center network provider type.
- Verify that you have scoped an existing NSX-T Data Center edge gateway to the data center group in which you want to create a routed network.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Data Center Group**, and select a group with an NSX-T Data Center network provider in which to create the network.
- 4 On the **Network Type** page, select **Routed** and click **Next**.

If there is only one available edge gateway scoped to the data center group, it is automatically assigned to the network.

- 5 If there is more than one NSX-T Data Center available to the data center group, select an edge gateway from the list, and click **Next**.
- 6 Enter a meaningful name for the network.
- 7 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.  
Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.
- 8 Enter a description of the organization VDC network.
- 9 Click **Next**.

- 10 (Optional) To reserve one or more IP addresses for assignment to virtual machines that require static IP addresses, configure the **Static IP Pools** for the network.
  - a Enter the IP address or range of IP addresses, and click **Add**.  
To add multiple static IP addresses or ranges, repeat this step.
  - b (Optional) To modify or remove IP addresses and ranges, click **Modify** or **Remove**.
- 11 (Optional) Configure the DNS settings.

Option	Action
Primary DNS	Enter the IP address for your primary DNS server.
Secondary DNS	Enter the IP address for your secondary DNS server.
DNS Suffix	Enter your DNS suffix. The DNS suffix is the DNS name without including the host name.

- 12 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Create a Data Center Group Network with an Imported NSX-T Logical Switch

**System administrators** can create an organization VDC network by importing a segment from an associated NSX-T Manager instance.

### Prerequisites

- Verify that you are a **system administrator**.
- Verify that you have created a data center group with an NSX-T Data Center network provider type.
- Verify that the provider virtual data center that backs the target virtual data center group is associated with an NSX-T Manager instance.
- Verify that you created at least one NSX-T logical switch that is not in use by other networks. For information about creating and configuring NSX-T logical switches, see the *NSX-T Data Center Administration Guide*.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Data Center Group**, and select a group with an NSX-T Data Center network provider in which to create the network.
- 4 On the **Network Type** page, select **Imported** and click **Next**.
- 5 From the list of available NSX-T logical switches, select the target switch and click **Next**.
- 6 Enter a meaningful name for the network.

- 7 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.  
Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.
- 8 Enter a description of the organization VDC network.
- 9 Click **Next**.
- 10 (Optional) To reserve one or more IP addresses for assignment to virtual machines that require static IP addresses, configure the **Static IP Pools** for the network.
  - a Enter the IP address or range of IP addresses, and click **Add**.  
To add multiple static IP addresses or ranges, repeat this step.
  - b (Optional) To modify or remove IP addresses and ranges, click **Modify** or **Remove**.
- 11 (Optional) Configure the DNS settings.

Option	Action
Primary DNS	Enter the IP address for your primary DNS server.
Secondary DNS	Enter the IP address for your secondary DNS server.
DNS Suffix	Enter your DNS suffix. The DNS suffix is the DNS name without including the host name.

- 12 On the **Ready to Complete** page, review your settings, and click **Finish**.

## Increase the Scope of an Organization VDC Network Backed by NSX-T Data Center

After you increase the scope of an organization VDC network to a data center group network, you can connect workloads from all data centers participating in the data center group.

### Prerequisites

- Verify that you are an **organization administrator** or that you have a role with an equivalent set of rights.
- Verify that you have created a data center group with an NSX-T Data Center network provider type.
- Verify that you created an organization VDC network backed by NSX-T Data Center.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the radio button next to the organization VDC network of which you want to increase the scope, and click **Increase Scope**.
- 3 Select a data center group from the list of data center groups, and click **OK** to confirm.

## Results

The scope of the network is increased to a data center group network. In the networks list, it is listed as scoped to the data center group that you selected.

## Decrease the Scope of a Data Center Group Network Backed by NSX-T Data Center

You can decrease the scope of a data center group network backed by NSX-T Data Center to an organization VDC network.

If you decrease the scope of a data center group network to a single organization VDC network, you provide network connectivity for workloads that belong only to the organization VDC.

### Prerequisites

- Verify that you are an **organization administrator** or that you have a role with an equivalent set of rights.
- Verify that you created a VDC network and you scoped it to a data center group with an NSX-T Data Center network provider type.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the radio button next to the data center group network of which you want to decrease the scope, and click **Decrease Scope**.
- 3 From the list of VDCs that are members of the group network, select the VDC to which you want to scope the network, and click **OK**.

## Results

The scope of the network is decreased to a single organization VDC network.

## Managing Egress Points for Data Center Groups with an NSX-T Data Center Network Provider Type

To route traffic in and out of a data center group network to an external network, you can configure an NSX-T Data Center edge gateway to be the egress point for a data center group.

When you configure an edge gateway to be the egress point for a data center group, you increase its scope to the data center group. The edge gateway becomes shared across all data centers that participate in the group. All routed networks that are attached to the edge gateway are attached to the data center group and scoped to it.

All edge gateway services remain part of the edge gateway functions. For more information, see [Managing NSX-T Data Center Edge Gateways](#).

If a VDC is a member of the data center group and if no workloads are attached to any of the routed networks that are not part of the targeted scope, you can remove an edge gateway from a data center group and scope it to a single VDC.

You can add an edge gateway to an isolated data center group network and convert it to a routed data center network. You can also remove the connection to an edge gateway from a data center group network, converting the routed network to an isolated data center group network.

## Add an NSX-T Data Center Edge Gateway to a Data Center Group

To configure an NSX-T Data Center edge gateway to be the egress point for a data center group, increase the scope of the edge gateway. The gateway then becomes shared across all data centers that participate in the group.

When you scope an edge gateway to a data center group, all routed networks that are attached to the edge gateway become attached to the data center group and scoped to it.

All new routed networks that you attach to the edge gateway belong to the data center group.

A routed network attached to an edge gateway which is scoped to a VDC can participate in a data center group only if the scope of the edge is increased to this data center group.

### Prerequisites

Verify that you have associated an existing NSX-T Data Center edge gateway with one of the VDCs that participate in the data center group.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click **Edge Gateway** and then click **Add Edge**.
- 4 Select one of the available edge gateways and click **Save**.

### Results

The scope of the edge gateway is increased to the data center group. The change of scope does not affect any existing underlying services or networks.

## Decrease the Scope of an NSX-T Data Center Edge Gateway to a VDC

You can decrease the scope of an NSX-T Data Center edge gateway to a specific VDC by removing the edge gateway from the data center group to which it is scoped.

When you decrease the scope of an edge gateway to a specific VDC, all security group objects that are in use by the edge gateway remain with it. Security groups that are used exclusively by the distributed firewall remain part of the VDC group.

### Prerequisites

- Verify that the VDC to which you want to decrease the scope of the edge gateway is a member of the data center group.

- Verify that there are no workloads attached to any routed networks that are not part of the targeted edge gateway scope.
- Verify that there are no security groups or IP sets in the data center group that are in use by both the edge gateway and the distributed firewall.

#### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.
- 3 Click **Edge Gateway** and then click **Remove Edge**.
- 4 Select a VDC to which to decrease the scope of the edge gateway and click **Save**.

## Managing Data Center Group Networking with NSX Data Center for vSphere

To create a network across multiple organization virtual data centers, you first group the virtual data centers, then create a VDC network that is scoped to the data center group.

VMware Cloud Director supports data center group networking for organization virtual data centers that are backed by NSX Data Center for vSphere with both an active and a stand-by egress point for a single network fault domain.

A data center group that is backed by NSX Data Center for vSphere can have either a common egress point configuration, an egress point configuration for each network fault domain, or a local group configuration.

#### Data center group

A data center group acts as a virtual data center group router that provides centralized networking administration, configuration for multiple egress points in multiple virtual data centers, and east-west traffic between all networks within the group. A data center group can contain between one and 16 virtual data centers that are configured to share multiple egress points. A data center group can have one of the following egress points configurations:

**Table 5-3. Egress Points Configuration Type for Data Center Groups Backed by NSX Data Center for vSphere**

Egress Points Configuration Type	Description
Common egress points configuration	<p>You can configure the data center group with one active egress point and one standby egress point. The two egress points are common to all participating virtual data centers across all network fault domains in the data center group.</p> <p>A data center group with this configuration can include data centers from up to four network fault domains.</p>
Egress points configuration per fault domain	<p>You can configure the data center group with one active egress point and one stand-by egress point for each network fault domain in the data center group.</p> <p>A data center group with this configuration can include data centers from up to four network fault domains.</p>
Local group configuration	<p>The organization virtual data centers in a local data center group are backed by a single vCenter Server instance. You can configure the local data center group with one active egress point and one standby egress point for a single network fault domain.</p>

An organization can have multiple data center groups. An organization virtual data center can participate in multiple data center groups.

The participating organization virtual data centers can belong to different VMware Cloud Director sites. See [Configure and Manage Multisite Deployments](#).

### Network Fault Domain

The network provider scope, typically representing the underlying vCenter Server instance with the associated NSX Manager.

### Egress point

An edge gateway that connects a data center group or network fault domain to the Internet. The edge gateway must belong to a virtual data center from the data center group. BGP routes are configured on the edge gateway representing the egress point and the universal router of the virtual data center group or network fault domain. Existing routes on the edge gateway are not affected.

### Stretched network

A layer 2 network that is stretched across all virtual data centers in a data center group. Can be IPv4 only.

## Managing Data Center Groups with NSX Data Center for vSphere Network Provider Type

After you create a data center group that is backed by NSX Data Center for vSphere, you can edit the network topology of a data center group. You can add and remove virtual data centers

from the group. You can swap, replace, and remove egress points. You can fix configuration failures by performing different synchronization tasks.

You cannot convert a common egress configuration to an egress configuration per fault domain or the reverse.

## Create and Configure a Data Center Group Backed by NSX Data Center for vSphere with a Common Egress Configuration

You can create and configure a virtual data center group backed by NSX Data Center for vSphere with a common egress configuration, where you set a pair of edge gateways that act as an active and stand-by egress points for all participating virtual data centers.

### Prerequisites

- This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.
- Your **system administrator** must enable the target virtual data centers for cross-virtual data center networking.

### Procedure

#### 1 [Create a Data Center Group Backed by NSX Data Center for vSphere with a Common Egress Configuration](#)

You can group between one and 16 virtual data centers in a data center group with a common egress configuration.

#### 2 [Add an Active Egress Point to a Data Center Group with NSX Data Center for vSphere Network Provider Type](#)

To connect your data center group to the Internet, you must add an active egress point to its network topology.

#### 3 [Add a Standby Egress Point to a Data Center Group with NSX Data Center for vSphere Network Provider Type](#)

In virtual data center groups with common egress configurations, you can add a secondary egress point, which acts as a standby egress point for fault tolerance scenarios.

## Create a Data Center Group Backed by NSX Data Center for vSphere with a Common Egress Configuration

You can group between one and 16 virtual data centers in a data center group with a common egress configuration.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click **New**.

- 3 On the **Starting VDC** page, select a VDC to start the VDC group.
- 4 Enter a name and, optionally, a description for the new data center group.
- 5 Select **Common Egress Points** and click **Next**.
- 6 On the **Participating VDCs** page, select additional data centers for the new data center group, and click **Next**.

The **Data Centers** page contains a list of the VDCs that the **system administrator** has enabled for cross-virtual data center networking.

- 7 Review the data center group details and click **Finish**.

### Results

The newly created virtual data center group is listed in the **Data Center Groups** view.

### Add an Active Egress Point to a Data Center Group with NSX Data Center for vSphere Network Provider Type

To connect your data center group to the Internet, you must add an active egress point to its network topology.

### Prerequisites

The **system administrator** created at least one edge gateway on any of the virtual data centers that are participating in the data center group.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 Click **Add egress point**.

The **Add Active Egress Point** page that opens provides a list of the edge gateways which belong to the participating virtual data centers.

- 4 Select the edge gateway that you want to act as an active egress point for this data center group, and click **Add**.

### Results

BGP routes are configured on the edge gateway representing the egress point and the universal router of the virtual data center group. Existing routes on the edge gateway are not affected.

The diagram of the network topology is updated with the newly added egress point. The traffic from the participating virtual data centers to the Internet is represented with a solid blue line.

## Add a Standby Egress Point to a Data Center Group with NSX Data Center for vSphere Network Provider Type

In virtual data center groups with common egress configurations, you can add a secondary egress point, which acts as a standby egress point for fault tolerance scenarios.

### Prerequisites

Apart from the edge gateway that acts as an active egress point, you must have at least one more edge gateway in any of the virtual data centers that are participating in the group.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 Click **Add stand-by egress point**.

The **Add Stand-by Egress Point** page opens providing a list of the unused edge gateways that belong to the participating virtual data centers. The edge gateway that is in use by the active egress point in this virtual data center group is not displayed.

- 4 Select the edge gateway that you want to act as a stand-by egress point for this data center group, and click **Add**.

### Results

BGP routes are configured on the edge gateway representing the egress point and the universal router of the network fault domain. The configuration does not affect the existing routes on the edge gateway.

The diagram of the network topology is updated with the newly added egress point. The traffic from the participating virtual data centers to the Internet in fault tolerance scenarios is represented with a dashed blue line.

## Create and Configure a Data Center Group Backed by NSX Data Center for vSphere with a Fault Domain Egress Configuration

You can create and configure a virtual data center group backed by NSX Data Center for vSphere with a fault domain egress configuration, where you configure an edge gateway that acts as an active egress points for each network fault domain in the group. Standby egresses cannot be created in a data center group with a fault domain egress configuration.

## Prerequisites

This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.

## Procedure

### 1 Create a Data Center Group Backed by NSX Data Center for vSphere with a Fault Domain Egress Configuration

You can group between 1 and 16 virtual data centers in a data center group backed by NSX Data Center for vSphere with a fault domain egress configuration.

### 2 Add an Egress Point for a Fault Domain

To connect the virtual data centers from a network fault domain in a data center group backed by NSX Data Center for vSphere to the Internet, you must add an egress point to this network fault domain. You can add an egress point to each network fault domain in the data center group. Stand-by egress points are not supported in a data center group with a fault domain egress configuration.

## Create a Data Center Group Backed by NSX Data Center for vSphere with a Fault Domain Egress Configuration

You can group between 1 and 16 virtual data centers in a data center group backed by NSX Data Center for vSphere with a fault domain egress configuration.

## Prerequisites

The **system administrator** enabled the target virtual data centers for cross-virtual data center networking.

## Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click **New**.

- 3 Enter a name and, optionally, a description for the new data center group.

- 4 Select **Egress Points per Fault Domain** and click **Next**.

- 5 On the **Participating VDCs** page, select additional data centers for the new data center group, and click **Next**.

The **Data Centers** page contains a list of the VDCs that the **system administrator** has enabled for cross-virtual data center networking.

- 6 Review the data center group details and click **Finish**.

## Results

The newly created virtual data center group is listed in the **Data Center Groups** view.

## Add an Egress Point for a Fault Domain

To connect the virtual data centers from a network fault domain in a data center group backed by NSX Data Center for vSphere to the Internet, you must add an egress point to this network fault domain. You can add an egress point to each network fault domain in the data center group. Stand-by egress points are not supported in a data center group with a fault domain egress configuration.

### Prerequisites

Apart from the edge gateways that are in use as egress points in this data center group, you must have at least one unused edge gateway in any of the participating virtual data centers.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.  
The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.
- 3 On the diagram of the network topology, click the target network fault domain.  
Network fault domains are represented with solid lines and their names at the bottom of the diagram.  
The selected fault domain is marked in blue.
- 4 Click **Add egress point**.  
The **Add Active Egress Point** page opens providing a list of the edge gateways that belong to the participating virtual data centers.
- 5 Select the edge gateway that you want to act as an egress point for this fault domain, and click **Add**.

### Results

BGP routes are configured on the edge gateway representing the egress point and the universal router of the network fault domain. Existing routes on the edge gateway are not affected.

The diagram of the network topology is updated with the newly added egress point. The traffic from the virtual data centers in the network fault domain to the Internet is represented with a continuous blue line.

## Create and Configure a Local Virtual Data Center Group with NSX Data Center for vSphere Network Provider Type

Starting with version 10.1, VMware Cloud Director supports data center groups backed by NSX Data Center for vSphere with both an active and a stand-by egress point for a single network fault domain.

The organization virtual data centers in a local group are backed by a single vCenter Server instance.

In a local data center group, you can set a pair of edge gateways - an active egress point and a stand-by egress point, to support high availability and disaster recovery scenarios within the same network fault domain.

### Prerequisites

This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.

### Procedure

- [1 Create a Local Data Center Group with NSX Data Center for vSphere Network Provider Type](#)  
You can group between 1 and 16 virtual data centers (VDCs) in a data center group backed by NSX Data Center for vSphere with a fault domain egress configuration.
- [2 Add an Active Egress Point for a Local Data Center Group with NSX Data Center for vSphere Network Provider Type](#)  
To connect the data centers from the local data center group backed by NSX Data Center for vSphere to the Internet, you must add an active egress point to the network fault domain.
- [3 Add a Stand-By Egress Point for a Local Data Center Group with NSX Data Center for vSphere Network Provider Type](#)  
In local data center groups configurations, you can add a secondary egress point, which acts as a stand-by egress point for fault tolerance scenarios.

### Create a Local Data Center Group with NSX Data Center for vSphere Network Provider Type

You can group between 1 and 16 virtual data centers (VDCs) in a data center group backed by NSX Data Center for vSphere with a fault domain egress configuration.

### Prerequisites

The **system administrator** enabled the target virtual data centers for cross-virtual data center networking.

### Procedure

- In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.

- 2 Click **New**.
- 3 On the **Starting VDC** page, select a VDC to start the VDC group.
- 4 Enter a name and, optionally, a description for the new data center group.
- 5 To create a group that contains only virtual data centers from a single network fault domain, toggle on the **Create Local Group** option.
- 6 Click **Next**.
- 7 On the **Participating VDCs** page, select additional data centers for the new data center group, and click **Next**.

The **Data Centers** page contains a list of the VDCs that the **system administrator** has enabled for cross-virtual data center networking.

- 8 Review the data center group details and click **Finish**.

### Results

The newly created virtual data center group appears in the **Data Center Groups** view.

### Add an Active Egress Point for a Local Data Center Group with NSX Data Center for vSphere Network Provider Type

To connect the data centers from the local data center group backed by NSX Data Center for vSphere to the Internet, you must add an active egress point to the network fault domain.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.  
The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.
- 3 Click **Add Egress Point**.
- 4 From the list of edge gateways that belong to the participating virtual data centers, select an edge gateway to act as an active egress point for the data center group, and click **Add**.

### Results

BGP routes are configured on the edge gateway representing the egress point and the universal router of the network fault domain. The configuration does not affect the existing routes on the edge gateway.

The newly added active egress point appears in the diagram of the network topology. A continuous blue line represents the traffic from the virtual data centers in the network fault domain to the Internet.

## What to do next

To allow for egress point fault tolerance, add a stand-by egress point for the local data center group.

### Add a Stand-By Egress Point for a Local Data Center Group with NSX Data Center for vSphere Network Provider Type

In local data center groups configurations, you can add a secondary egress point, which acts as a stand-by egress point for fault tolerance scenarios.

#### Prerequisites

Apart from the edge gateway that acts as an active egress point, you must have at least one more edge gateway in any of the virtual data centers that are participating in the local data center group.

#### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 Click **Add stand-by egress point**.

The **Add Stand-by Egress Point** page opens providing a list of the unused edge gateways that belong to the participating virtual data centers. The edge gateway that is in use by the active egress point in this virtual data center group appears dimmed.

- 4 Select the edge gateway that you want to act as a stand-by egress point for this data center group, and click **Add**.

#### Results

BGP routes are configured on the edge gateway representing the egress point and the universal router of the network fault domain. The configuration does not affect the existing routes on the edge gateway.

The newly added egress point appears in the network topology diagram. A dashed blue line represents the traffic from the participating virtual data centers to the Internet in fault tolerance scenarios.

### View a Data Center Group with NSX Data Center for vSphere Network Provider Type

You can view the data center groups in your organization and details about their current configuration.

## Prerequisites

This operation requires the **System Administrator** role or a role with the **VDC Group: View VDC Group** right published to the organization.

## Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

## Add a Virtual Data Center to a Data Center Group with NSX Data Center for vSphere Network Provider Type

You can add a virtual data center to a data center group, as a result stretching the existing networks to the new virtual data center.

## Prerequisites

- This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.
- The data center group contains less than four virtual data centers.

## Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 Click **Add Data Center**.

- 4 On the **Data Centers** page, select the data center that you want to add to the data center group and click **Finish**.

The **Data Centers** page contains a list virtual data centers that are enabled for cross-virtual data center networking by the system administrator.

---

**Note** A data center group must contain up to four virtual data centers.

---

## Remove a Virtual Data Center from a Data Center Group with NSX Data Center for vSphere Network Provider Type

You can remove a virtual data center from a data center group, as a result unstretching the existing networks from this virtual data center.

### Prerequisites

- This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.
- The data center group must contain at least three virtual data centers.
- The virtual data center that you want to remove must not provide an egress point to the data center group.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.  
The list of data center groups appears.
- 2 Click the target data center group.  
The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.
- 3 In the upper right corner of the card of the target virtual data center, click the three dots, and click **Remove**.
- 4 To confirm, click **Remove**.

### Results

The virtual data center is removed from the network topology diagram of the data center group.

## Synchronize a Data Center Group with NSX Data Center for vSphere Network Provider Type

To reapply the data center group network configurations and ensure that all participating virtual data centers are active, you can synchronize the data center group.

---

**Note** During the data center group synchronization process, the data center group becomes unavailable for a few seconds, because the universal router synchronizes in NSX.

---

### Prerequisites

This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 Click **Sync data center group**.

- 4 To confirm, click **OK**.

### Swap the Egress Points in a Data Center Group with NSX Data Center for vSphere Network Provider Type and a Common Egress Configuration

After you configure an active and stand-by egress points in a data center group with a common egress configuration, you can swap the roles of the egress points. The active egress point can become a stand-by egress point and the reverse.

### Prerequisites

This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 Click **Swap egress points**.

- 4 To confirm, click **OK**.

### Results

The diagram of the network topology is updated with the new traffic routes. The traffic to the Internet is now redirected to the new active egress point.

### Replace the Edge Gateway of an Egress Point of a Data Center Group with NSX Data Center for vSphere Network Provider Type

You can replace the edge gateway that represents an active or standby egress point in a data center group.

### Prerequisites

- This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.
- The new edge gateway must not be in use by other egress points in the data center group.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 If you are replacing an egress point from a network fault domain configuration, on the network topology diagram, select the network fault domain of the target egress point.

Network fault domains are represented with solid lines and domain names at the bottom of the diagram.

The selected network fault domain is marked in blue.

- 4 In the upper right corner of the card of the target egress point, click the three dots, and click **Replace**.

The **Replace Egress Point** page opens providing a list of the edge gateways that belong to the participating virtual data centers.

- 5 Select the new edge gateway and click **Replace**.

### Results

BGP routes are removed from the old edge gateway and configured on the new edge gateway representing the egress point and the universal router of the virtual data center group.

The network topology diagram is updated with the name of the new edge gateway.

## Remove an Egress Point from a Data Center Group with NSX Data Center for vSphere Network Provider Type

To disconnect a data center group or network fault domain from the Internet, you can remove its egress point.

### Prerequisites

- This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.
- If you want to remove an active egress point that is paired with a stand-by egress point, you must swap the egress points or remove the stand-by egress point.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 If you are removing an egress point from a network fault domain configuration, on the network topology diagram, select the network fault domain of the target egress point.

Network fault domains are represented with solid lines and domain names at the bottom of the diagram.

The selected network fault domain is marked in blue.

- 4 In the upper right corner of the card of the target egress point, click the three dots, and click **Delete**.

- 5 To confirm, click **OK**.

### Results

BGP routes are removed from the edge gateway representing the egress point if it is not in use by other universal routers.

The egress point is removed from the network topology diagram.

## Synchronize Routes and Egress Points of a Data Center Group with NSX Data Center for vSphere Network Provider Type

You can reapply the dynamic routing configuration to a data center group or network fault domain and its associated egress points by synchronizing the routes. You can ensure that an egress point is properly connected to the data center group by synchronizing the egress point.

### Prerequisites

- This operation requires the **System Administrator** role or a role with the **VDC Group: Configure VDC Group** right published to the organization.
- You configured an egress point for the target data center group or network fault domain.

### Procedure

- 1 In the top navigation bar, click **Networking** and then click the **Data Center Groups** tab.

The list of data center groups appears.

- 2 Click the target data center group.

The **Network Topology** view for this data center group opens. The diagram of the current network topology displays the participating VDCs with their network fault domains, the egress points, if configured, and the traffic routes.

- 3 If you are synchronizing a network fault domain in a data center group, on the network topology diagram, select the target network fault domain.

Network fault domains are represented with solid lines and domain names at the bottom of the diagram.

The selected network fault domain is marked in blue.

- 4 To reapply the dynamic routing configuration to the group or network fault domain and its associated egress points, click **Sync routes**, and click **OK**.
- 5 To synchronize an egress point with its data center group, in the upper right corner of the card of the target egress point, click the three dots, click **Sync**, and click **OK**.

## Managing Data Center Group Networks Backed by NSX Data Center for vSphere

After you create and configure a data center group, you can create and manage VDC group layer 2 networks spanning the participating virtual data centers.

### Add a VDC Group Network Backed by NSX Data Center for vSphere

You can create a VDC group network across all virtual data centers that are participating in a data center group.

You can add only an IPv4 data center group network backed by NSX Data Center for vSphere.

#### Prerequisites

This operation requires the predefined **Organization Administrator** role or a role with the **Organization VDC Network: Edit Properties** right.

#### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 On the **Networks** tab, click **New**.
- 3 On the **Scope** page, select **Data Center Group**, and select a data center group backed by NSX Data Center for vSphere in which to create the network and click **Next**.
- 4 Enter a meaningful name for the network.
- 5 Enter the Classless Inter-Domain Routing (CIDR) settings for the network.  
Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.
- 6 Enter a description of the organization VDC network.

- 7 Click **Next**.
- 8 Review the settings and click **Finish**.

### Results

You can see the newly created data center group network in the list of network for the organization.

Its network type is listed as Cross-VDC.

An organization virtual data center network of cross-VDC routing type is created for each participating virtual data center. You can see the VDC group networks of the participating virtual data centers by clicking on the card of a participating virtual data center and then clicking **Networks**. If a virtual machine or vApp connects to such an organization virtual data center network, this virtual machine or vApp connects to the VDC group network.

### What to do next

For each corresponding cross-VDC organization virtual data center network, you can assign static IP addresses and IP pools. See [Add IP Addresses to an Organization Virtual Data Center Network IP Pool](#).

For DNS and DHCP configurations for virtual machines attached to a VDC group network, you can use the VMware Cloud Director OpenAPI. To examine the VMware Cloud Director OpenAPI documentation, go to [https://Cloud\\_Director\\_IP\\_address\\_or\\_host\\_name/docs](https://Cloud_Director_IP_address_or_host_name/docs). To view code samples and test VMware Cloud Director OpenAPI calls, go to [https://Cloud\\_Director\\_IP\\_address\\_or\\_host\\_name/api-explorer?scope=organization\\_name](https://Cloud_Director_IP_address_or_host_name/api-explorer?scope=organization_name).

## View or Edit a Data Center Group Network Backed by NSX Data Center for vSphere

You can view the name, the description, and the CIDR settings of a data center group network backed by NSX Data Center for vSphere. You can edit only the name and description of a data center group network backed by NSX Data Center for vSphere.

For information about editing the static IP pool allocation for a data center group network at a virtual data center level, see [Add IP Addresses to an Organization Virtual Data Center Network IP Pool](#).

### Prerequisites

Verify that you are assigned the predefined **Organization Administrator** role or a role that includes the **Organization VDC Network: View Properties** and the **Organization VDC Network: Edit Properties** right.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 Click the target network to view its details.
- 3 To edit the name and the description of the networks, click **Edit**.

- 4 Edit the network details and click **Save**.

## Synchronize a Data Center Group Network Backed by NSX Data Center for vSphere

To ensure that all participating virtual data centers can access their data center group network backed by NSX Data Center for vSphere, you can synchronize the data center group network.

### Prerequisites

This operation requires the predefined **Organization Administrator** role or a role with the **Organization VDC Network: Edit Properties** right.

### Procedure

- 1 In the top navigation bar, click **Networking**.
- 2 In the networks tab, select the radio button next to the name of the target network, and click **Sync**.
- 3 To confirm, click **OK**.

## Managing NSX Data Center for vSphere Edge Gateway Services

VMware Cloud Director provides advanced networking capabilities powered by the NSX Data Center for vSphere network virtualization software that offer enhanced security controls, routing, and network scaling capabilities in a cloud environment.

Using these networking capabilities, you can achieve unprecedented security and isolation in your organization virtual data center. These capabilities deliver the following benefits:

- **Dynamic routing.** The NSX Data Center for vSphere capabilities in your VMware Cloud Director environment support routing protocols such as Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF) to simplify network integration between systems, to provide redundancy and continuity in a cloud-hosted application deployment.
- **Fine-grained network security and isolation.** The NSX Data Center for vSphere capabilities in your VMware Cloud Director environment support the use of object-based rule definitions to provide stateful network traffic isolation without requiring multiple virtual networks. This zero-trust security model prevents intruders from gaining full network access if an application or virtual machine is compromised. Network configuration is simplified by using the same network security policies to protect applications wherever they are physically located in the VMware Cloud Director environment and to extend your zero-trust security model for portable security no matter where an application is deployed.
- **Additional capabilities provided by NSX Data Center for vSphere** are enhanced VPN support for point-to-site (IPsec VPN) and user (SSL VPN-Plus) connectivity, enhanced load balancing for HTTPS, and expanded network scalability.

You can configure two types of firewalls: the edge gateway firewall and the distributed firewall. For more information about the differences between these firewalls, see [Tenant Firewall Configuration with NSX Data Center for vSphere](#).

You access these advanced networking capabilities using the VMware Cloud Director Tenant Portal or the VMware Cloud Director Service Provider Admin Portal. The edge gateway must first be converted to an advanced edge gateway. See [Convert an NSX Data Center for vSphere Edge Gateway to an Advanced Edge Gateway](#).

---

**Important** IPv6 edge gateways support limited services. IPv6 edge gateways support edge firewalls, distribute firewalls, and static routing.

---

## Getting Started with VMware Cloud Director Advanced Networking with NSX Data Center for vSphere

You use the VMware Cloud Director Advanced Networking to perform management tasks on an organization in a VMware Cloud Director system. You can manage distributed firewalls and other advanced networking capabilities that are provided by NSX Data Center for vSphere and made available to an organization by a VMware Cloud Director system administrator.

The typical users of the advanced networking provided by NSX Data Center for vSphere are:

- VMware Cloud Director **system administrators**, who might use the tenant portal to configure the distributed firewall and other advanced networking capabilities for an organization.
- **Organization administrators**, who use the tenant portal to manage the distributed firewall and other advanced networking capabilities that the **system administrator** has made available to that organization.

## Tenant Firewall Configuration with NSX Data Center for vSphere

Using the tenant portal, you can configure the firewall capabilities provided by NSX Data Center for vSphere in your VMware Cloud Director organization virtual data center. You can create firewall rules for distributed firewalls to provide security between virtual machines in an organization virtual data center and firewall rules to apply to an edge gateway firewall to protect the virtual machines in an organization virtual data center from outside network traffic.

---

**Note** The tenant portal provides the ability to configure both edge gateway firewalls and distributed firewalls.

---

The NSX Data Center for vSphere logical firewall technology consists of two components to address different deployment use cases. The edge gateway firewall focuses on North-South traffic enforcement while the distributed firewall focuses on East-West access controls.

### Key Differences Between Edge Gateway Firewalls and Distributed Firewalls

An edge gateway firewall monitors North-South traffic to provide perimeter security functionality including firewall, Network Address Translation (NAT) as well as site-to-site IPsec and SSL VPN functionality.

A distributed firewall provides the capability to isolate and secure each virtual machine and application down to the layer 2 (L2) level. Configuring distributed firewalls effectively quarantines any external or internal network security compromise, isolating East-West traffic between virtual machines on the same network segment. Security policies are centrally managed, inheritable, and nestable, so networking and security administrators can manage them at scale. Additionally, once deployed, defined security policies follow the virtual machines or applications when they move between different virtual data centers.

## About Firewall Rules

As described in the relevant product documentation, in NSX Data Center for vSphere, the firewall rules defined on the centralized level are referred to as pre rules. You can also add rules at an individual edge gateway level, and those rules are referred to as local rules.

Each traffic session is checked against the top rule in the firewall table before moving down the subsequent rules in the table. The first rule in the table that matches the traffic parameters is enforced. Rules are displayed in the following order:

- 1 User-defined pre rules have the highest priority, and are enforced in top-to-bottom ordering with a per-virtual NIC level precedence.
- 2 Auto-plumbed rules (rules that enable control traffic to flow for edge gateway services).
- 3 Local rules defined at an edge gateway level.
- 4 Default distributed firewall rule

For more information about how the NSX Data Center for vSphere software enforces firewall rules, see *Change the Order of a Firewall Rule* in the NSX Data Center for vSphere documentation.

## NSX Data Center for vSphere Edge Gateway Firewall

The firewall for the edge gateway helps you meet key perimeter security requirements, such as building DMZs based on IP/VLAN constructs, tenant-to-tenant isolation in multi-tenant virtual data centers, Network Address Translation (NAT), partner (extranet) VPNs, and user-based SSL VPNs.

The edge gateway firewall capability in the VMware Cloud Director environment is provided by NSX Data Center for vSphere. In NSX Data Center for vSphere, this firewall capability is also referred to as the edge firewall. The edge gateway firewall monitors North-South traffic to provide perimeter security functionality including firewall, Network Address Translation (NAT) as well as site-to-site IPSec and SSL VPN functionality.

For more detailed information about the capabilities provided by the edge gateway firewall of NSX Data Center for vSphere, see the NSX Data Center for vSphere documentation.

## Managing an NSX Data Center for vSphere Edge Gateway Firewall

To protect traffic to and from an edge gateway, you can create and manage firewall rules on that edge gateway.

For information about protecting traffic traveling between virtual machines in an organization virtual data center, see [Managing NSX Data Center for vSphere Distributed Firewall Rules Using the Tenant Portal](#).

Rules created on the distributed firewall screen that have an advanced edge gateway specified in their Applied To column are not displayed in the Firewall screen for that advanced edge gateway .

The edge gateway firewall rules for an edge gateway are displayed in the **Firewall** screen and are enforced in the following order:

- 1 Internal rules, also known as auto-plumbed rules. These internal rules enable control traffic to flow for edge gateway services.
- 2 User-defined rules.
- 3 Default rule.

The default rule settings apply to traffic that does not match any of the user-defined firewall rules. The default rule is displayed at the bottom of the rules on the Firewall screen.

In the tenant portal, use the **Enable** toggle on the Firewall Rules screen of the edge gateway to activate or deactivate an edge gateway firewall.

### Convert an NSX Data Center for vSphereEdge Gateway to an Advanced Edge Gateway

To work with an NSX Data Center for vSphere edge gateway in the tenant portal, you need to convert it to an advanced edge gateway. Once you convert it to an advanced edge gateway, you can use the tenant portal to configure the static and dynamic routing capabilities that are provided by NSX Data Center for vSphere for those advanced edge gateways.

#### Prerequisites

You have an existing edge gateway.

#### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Select the edge gateway to edit.
- 3 Click **Convert to Advanced**.

#### Results

Your edge gateway is converted to an advanced edge gateway.

#### What to do next

Once you have converted to an advanced edge gateway, you can configure settings by selecting the gateway and clicking **Services**.

## Add an NSX Data Center for vSphere Edge Gateway Firewall Rule

You use the edge gateway **Firewall** tab to add firewall rules for that edge gateway. You can add multiple NSX Edge interfaces and multiple IP address groups as the source and destination for these firewall rules.

Specifying **internal** for a source or a destination of a rule indicates traffic for all subnets on the port groups connected to the NSX edge gateway. If you select **internal** as the source, the rule is automatically updated when additional internal interfaces are configured on the NSX gateway.

---

**Note** Edge gateway firewall rules on internal interfaces do not work when the edge gateway is configured for dynamic routing.

---

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 If the **Firewall Rules** screen is not already visible, click the **Firewall** tab.
- 3 To add a rule below an existing rule in the firewall rules table, click in the existing row and then click the **Create** button.

A row for the new rule is added below the selected rule, and is assigned any destination, any service, and the **Allow** action by default. When the system-defined default rule is the only rule in the firewall table, the new rule is added above the default rule.

- 4 Click in the **Name** cell and type in a name.
- 5 Click in the **Source** cell and use the now visible icons to select a source to add to the rule:

Option	Description
Click the IP icon	Type the source value you want to use. Valid values are an IP address, CIDR, an IP range, or the keyword <b>any</b> . The edge gateway firewall supports both IPv4 and IPv6 formats.
Click the + icon	<p>Use the + icon to specify the source as an object other than a specific IP address:</p> <ul style="list-style-type: none"> <li>■ Use the <b>Select objects</b> window to add objects that match your selections and click <b>Keep</b> to add them to the rule.</li> <li>■ To exclude a source from the rule, add it to this rule using the <b>Select objects</b> window and then select the toggle exclusion icon to exclude that source from this rule.</li> </ul> <p>When the toggle exclusion is selected on the source, the rule is applied to traffic coming from all sources except for the source you excluded. When the toggle exclusion is not selected, the rule applies to traffic coming from the source you specified in the <b>Select objects</b> window</p>

- 6 Click in the **Destination** cell and perform one of the following options:

Option	Description
Click the IP icon	Type the destination value you want to use. Valid values are an IP address, CIDR, an IP range, or the keyword <b>any</b> . The edge gateway firewall supports both IPv4 and IPv6 formats.
Click the + icon	<p>Use the + icon to specify the source as an object other than a specific IP address:</p> <ul style="list-style-type: none"> <li>■ Use the <b>Select objects</b> window to add objects that match your selections and click <b>Keep</b> to add them to the rule.</li> <li>■ To exclude a source from the rule, add it to this rule using the <b>Select objects</b> window and then select the toggle exclusion icon to exclude that source from this rule.</li> </ul> <p>When the toggle exclusion is selected on the source, the rule is applied to traffic coming from all sources except for the source you excluded. When the toggle exclusion is not selected, the rule applies to traffic coming from the source you specified in the <b>Select objects</b> window</p>

- 7 Click in the **Service** cell of the new rule and click the + icon to specify the service as a port-protocol combination:
- a Select the service protocol.
  - b Type the port numbers for the source and destination ports, or specify **any**.
  - c Click **Keep**.
- 8 In the **Action** cell of the new rule, configure the action for the rule.

Option	Description
<b>Accept</b>	Allows traffic from or to the specified sources, destinations, and services.
<b>Deny</b>	Blocks traffic from or to the specified sources, destinations, and services.

- 9 Click **Save changes**.

The save operation can take a minute to complete.

### Modify NSX Data Center for vSphere Edge Gateway Firewall Rules

You can edit and delete only the user-defined firewall rules that were added to an edge gateway. You cannot edit or delete an auto-generated rule or a default rule, except for changing the action setting of the default rule. You can change the priority order of user-defined rules.

For details about the available settings for the various cells of a rule, see [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#).

## Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Firewall** tab.
- 3 Manage the firewall rules.
  - Deactivate a rule by clicking the green check mark in its **No.** cell. The green check mark turns to a red deactivated icon. If the rule is deactivated and you want to activate the rule, click the red deactivated icon.
  - Edit a rule name by double-clicking in its **Name** cell and typing the new name.
  - Modify the settings for a rule, such as the source or action settings, by selecting the appropriate cell and using the displayed controls.
  - Delete a rule by selecting it and clicking the **Delete** button located above the rules table.
  - Hide system-generated rules by using the **Show only user-defined rules** toggle.
  - Move a rule up or down in the rules table by selecting the rule and clicking the up and down arrow buttons located above the rules table.
- 4 Click **Save changes**.

## NSX Data Center for vSphere Distributed Firewall

The distributed firewall allows you to segment organization virtual data center entities, such as virtual machines, based on virtual machine names and attributes.

VMware Cloud Director supports distributed firewall services on organization virtual data centers that are backed by NSX Data Center for vSphere. As described in the NSX Data Center for vSphere documentation, this distributed firewall is a hypervisor kernel-embedded firewall that provides visibility and control for virtualized workloads and networks. You can create access control policies based on objects like virtual machine names and on network constructs like IP addresses or IP set addresses. Firewall rules are enforced at the vNIC level of each virtual machine to provide consistent access control even when the virtual machine is moved to a new ESXi host by vSphere vMotion. This distributed firewall supports a micro-segmentation security model where East-West traffic can be inspected at near line rate processing.

As described in the NSX Data Center for vSphere documentation, for layer 2 (L2) packets, the distributed firewall creates a cache for performance boost. Layer 3 (L3) packets are processed in the following sequence:

- 1 All packets are checked for an existing state.
- 2 When a state match is found, the packets are processed.

- 3 When a state match is not found, the packets are processed through the rules until a match is found.
  - For TCP packets, a state is set only for packets with a SYN flag. However, rules that do not specify a protocol (service ANY), can match TCP packets with any combination of flags.
  - For UDP packets, 5-tuple details are extracted from the packet. When a state does not exist in the state table, a new state is created using the extracted 5-tuple details. Subsequently received packets are matched against the state that was just created.
  - For ICMP packets, ICMP type, code, and packet direction are used to create a state.

The distributed firewall can help in creating identity-based rules as well. Administrators can enforce access control based on the user's group membership as defined in the enterprise Active Directory (AD). Some use cases for when you might use identity-based firewall rules are:

- Users accessing virtual applications using a laptop or mobile device where AD is used for user authentication
- Users accessing virtual applications using VDI infrastructure where the virtual machines are Microsoft Windows-based

For more detailed information about the capabilities provided by the distributed firewall, see the NSX Data Center for vSphere documentation.

## Enable the Distributed Firewall on an Organization Virtual Data Center Backed by NSX Data Center for vSphere

Before you can use the tenant portal to work with the distributed firewall capabilities provided by NSX Data Center for vSphere on an organization virtual data center, the distributed firewall must be enabled for that organization virtual data center. A VMware Cloud Director system administrator or a user granted the **org\_vdc\_distributed\_firewall\_enable** right can enable the distributed firewall on an organization virtual data center.

You use the Distributed Firewall screen in the tenant portal to enable the distributed firewall for an organization virtual data center.

### Prerequisites

Verify that the organization to which the organization virtual data center belongs has the following rights assigned to it:

- Organization vDC Distributed Firewall: Enable/Disable
- Organization vDC Distributed Firewall: Configure Rules
- Organization vDC Distributed Firewall: View Rules

The VMware Cloud Director **system administrator** assigns rights to an organization. The Organization vDC Distributed Firewall: Enable/Disable right is required for activating the distributed firewall using the user interface in the tenant portal. The Organization vDC Distributed Firewall: View Rules right is required for viewing the firewall rules in the tenant portal and the Organization vDC Distributed Firewall: Configure Rules right is required for configuring the firewall rules using the tenant portal.

Verify that you have an assigned role that grants you the right named Organization vDC Distributed Firewall: Enable/Disable. Of the pre-defined roles in a VMware Cloud Director system, only the System Administrator role has that right by default.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.
- 2 Select the organization virtual data center for which you want to configure distributed firewall rules.
- 3 Click **Configure Services**.
- 4 Enable distributed firewall on the **Distributed Firewall** tab.

### What to do next

For a description of the default distributed firewall rule, see [Managing NSX Data Center for vSphere Distributed Firewall Rules Using the Tenant Portal](#).

## Managing NSX Data Center for vSphere Distributed Firewall Rules Using the Tenant Portal

As described in the NSX Data Center for vSphere documentation, default firewall settings apply to traffic that does not match any of the user-defined firewall rules. In the VMware Cloud Director Tenant Portal, the default distributed firewall rule is labeled Default Allow Rule.

The distributed firewall capability must be enabled on an organization virtual data center before you can manage the distributed firewall settings using the VMware Cloud Director Tenant Portal.

The default distributed firewall rule is configured to allow all layer 3 and layer 2 traffic to pass through the organization virtual data center. This setting is indicated by the Allow set in the Action column in the user interface. The default rule is always at the bottom of the rules table.

---

**Important** You cannot delete or modify the default distributed firewall rules.

---

### Add a Distributed Firewall Rule

You first add a distributed firewall rule to the scope of the organization virtual data center. Then you can narrow down the scope at which you want to apply the rule. The distributed firewall allows you to add multiple objects at the source and destination levels for each rule, which helps reduce the total number of firewall rules to be added.

For information about the predefined services and service groups that you can use in a rule, see [View Services Available for Firewall Rules](#) and [View Service Groups Available for Firewall Rules](#).

### Prerequisites

- [Enable the Distributed Firewall on an Organization Virtual Data Center Backed by NSX Data Center for vSphere](#)
- If you want to use an IP set as a source or destination in a rule, [Create an IP Set for Use in Firewall Rules and DHCP Relay Configuration](#).
- If you want to use a MAC set as a source or destination in a rule, [Create a MAC Set for Use in Firewall Rules](#).
- If you want to use a Security group as a source or destination in a rule, [Create a Security Group](#).

### Procedure

1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.

2 Select the security services VDC network for which you want to modify firewall rules, and click **Configure Services**.

The Security Services screen displays.

3 Select the type of rule you want to create. You have the option to create a general rule or an Ethernet rule.

Layer 3 (L3) rules are configured on the **General** tab. Layer 2 (L2) rules are configured on the **Ethernet** tab.

4 To add a rule below an existing rule in the firewall table, click in the existing row and then click the **Create** () button.

A row for the new rule is added below the selected rule, and is assigned any destination, any service, and the **Allow** action by default. When the system-defined Default Allow rule is the only rule in the firewall table, the new rule is added above the default rule.

5 Click in the **Name** cell and type in a name.

## 6 Click in the **Source** cell and use the now visible icons to select a source to add to the rule:

Action	Description
Click the IP icon	Applicable for rules defined on the <b>General</b> tab. Type the source value you want to use. Valid values are an IP address, CIDR, an IP range, or the keyword <b>any</b> . The distributed firewall supports IPv4 format only.
Click the + icon	Use the + icon to specify the source as an object other than a specific IP address: <ul style="list-style-type: none"> <li>■ Use the <b>Select objects</b> window to add objects that match your selections and click <b>Keep</b> to add them to the rule.</li> <li>■ To exclude a source from the rule, add it to this rule using the <b>Select objects</b> window and then select the toggle exclusion icon to exclude that source from this rule.</li> </ul> <p>When the toggle exclusion is selected on the source, the rule is applied to traffic coming from all sources except for the source you excluded. When the toggle exclusion is not selected, the rule applies to traffic coming from the source you specified in the <b>Select objects</b> window</p>

## 7 Click in the **Destination** cell and perform one of the following actions:

Action	Description
Click the IP icon	Applicable for rules defined on the <b>General</b> tab. Type the destination value you want to use. Valid values are an IP address, CIDR, an IP range, or the keyword <b>any</b> . The distributed firewall supports IPv4 format only.
Click the + icon	Use the + icon to specify the source as an object other than a specific IP address: <ul style="list-style-type: none"> <li>■ Use the <b>Select objects</b> window to add objects that match your selections and click <b>Keep</b> to add them to the rule.</li> <li>■ To exclude a source from the rule, add it to this rule using the <b>Select objects</b> window and then select the toggle exclusion icon to exclude that source from this rule.</li> </ul> <p>When the toggle exclusion is selected on the source, the rule is applied to traffic coming from all sources except for the source you excluded. When the toggle exclusion is not selected, the rule applies to traffic coming from the source you specified in the <b>Select objects</b> window</p>

- 8 Click in the **Service** cell of the new rule and perform one of the following actions:

Action	Description
Click the IP icon	To specify the service as a port–protocol combination: <ol style="list-style-type: none"> <li>Select the service protocol.</li> <li>Type the port numbers for the source and destination ports, or specify <b>any</b>, and click <b>Keep</b>.</li> </ol>
Click the + icon	To select a pre-defined service or service group, or define a new one: <ol style="list-style-type: none"> <li>Select one or more objects and add them to the filter.</li> <li>Click <b>Keep</b>.</li> </ol>

- 9 In the **Action** cell of the new rule, configure the action for the rule.

Option	Description
Allow	Allows traffic from or to the specified sources, destinations, and services.
Deny	Blocks traffic from or to the specified sources, destinations, and services.

- 10 In the **Direction** cell of the new rule, select whether the rule applies to incoming traffic, outgoing traffic, or both.
- 11 If this is a rule on the **General** tab, in the **Packet Type** cell of the new rule, select a packet type of **Any**, **IPV4**, or **IPV6**.
- 12 Select the **Applied To** cell, and use the + icon to define the object scope to which this rule is applicable.

When the rule contains virtual machines in the **Source** and **Destination** cells, you must add both the source and destination virtual machines to the rule's **Applied To** for the rule to work correctly.

**Important** IP address groups (IP sets), MAC address groups (MAC sets), and security groups containing either IP sets or MAC sets are not valid input parameters.

- 13 Click **Save Changes**.

### Edit a Distributed Firewall Rule

In a VMware Cloud Director environment, to modify an existing distributed firewall rule of an organization virtual data center, use the **Distributed Firewall** screen.

For details about the available settings for the various cells of a rule, see [Add a Distributed Firewall Rule](#).

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.

- 2 Select the security services VDC network for which you want to modify firewall rules, and click **Configure Services**.

The Security Services screen displays.

- 3 Perform any of the following actions to manage the distributed firewall rules:
  - Deactivate a rule by clicking the green check mark in its **No.** cell.  
The green check mark turns to a red deactivated icon. If the rule is deactivated and you want to activate the rule, click the red deactivated icon.
  - Edit a rule name by double-clicking in its **Name** cell and typing the new name.
  - Modify the settings for a rule, such as the source or action settings, by selecting the appropriate cell and using the displayed controls.
  - Delete a rule by selecting it and clicking the **Delete** button located above the rules table.
  - Move a rule up or down in the rules table by selecting the rule and clicking the up and down arrow buttons located above the rules table.
- 4 Click **Save Changes**.

## Managing NSX Data Center for vSphere Edge Gateway DHCP

You configure your edge gateways to provide Dynamic Host Configuration Protocol (DHCP) services to virtual machines connected to the associated organization virtual data center networks.

As described in the [NSX documentation](#), an NSX edge gateway capabilities include IP address pooling, one-to-one static IP address allocation, and external DNS server configuration. Static IP address binding is based on the managed object ID and interface ID of the requesting client virtual machine.

The DHCP service for an NSX edge gateway:

- Listens on the internal interface of the edge gateway for DHCP discovery.
- Uses the IP address of the internal interface of the edge gateway as the default gateway address for all clients.
- Uses the broadcast and subnet mask values of the internal interface for the container network.

In the following situations, you need to restart the DHCP service on the client virtual machines that have the DHCP-assigned IP addresses:

- You changed or deleted a DHCP pool, default gateway, or DNS server.

- You changed the internal IP address of the edge gateway instance.

**Note** If the DNS settings on a edge gateway which has DHCP activated are changed, the edge gateway might stop providing DHCP services. If this situation occurs, use the **DHCP Service Status** toggle on the DHCP Pools screen to deactivate and then reactivate DHCP on that edge gateway. See [Add a DHCP IP Pool](#).

## Add a DHCP IP Pool

You can configure the IP pools needed for a DHCP service of an NSX Data Center for vSphere edge gateway. DHCP automates IP address assignment to virtual machines connected to organization virtual data center networks.

As described in the *NSX Administration* documentation, the DHCP service requires a pool of IP addresses. An IP pool is a sequential range of IP addresses within the network. Virtual machines protected by the edge gateway that do not have an address binding are allocated an IP address from this pool. IP pool ranges cannot intersect one another, thus one IP address can belong to only one IP pool.

**Note** At least one DHCP IP pool must be configured to have the DHCP service status turned on.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **DHCP > Pools** .
- 3 If DHCP service is not currently enabled, turn on the **DHCP Service Status** toggle.

**Note** Add at least one DHCP IP pool before saving changes after turning on the **DHCP Service Status** toggle. If no DHCP IP pools are listed on the screen and you turn on the **DHCP Service Status** toggle and save the changes, the screen displays with the toggle turned off.

- 4 Under DHCP Pools, click the **Create** () button, specify the details for the DHCP pool, and click **Keep**.

Option	Description
IP Range	Type in a range of IP addresses.
Domain Name	Domain name of the DNS server.
Auto Configure DNS	Turn on this toggle to use the DNS service configuration for this IP pool DNS binding. If enabled, the <b>Primary Name Server</b> and <b>Secondary Name Server</b> are set to <b>Auto</b> .

Option	Description
<b>Primary Name Server</b>	When you do not enable <b>Auto Configure DNS</b> , type your primary DNS server IP address of your primary DNS server. This IP address is used for hostname-to-IP address resolution.
<b>Secondary Name Server</b>	When you do not enable <b>Auto Configure DNS</b> , type your secondary DNS server IP address. This IP address is used for hostname-to-IP address resolution.
<b>Default Gateway</b>	Type the default gateway address. When you do not specify the default gateway IP address, the internal interface of the edge gateway instance is taken as the default gateway.
<b>Subnet Mask</b>	Type the subnet mask of the edge gateway interface.
<b>Lease Never Expires</b>	Enable this toggle to keep the IP addresses that are assigned out of this pool bound to their assigned virtual machines forever. When you select this option, <b>Lease Time</b> is set to infinite.
<b>Lease Time (Seconds)</b>	Length of time (in seconds) that the DHCP-assigned IP addresses are leased to the clients. The default lease time is one day (86400 seconds).
	<b>Note</b> You cannot specify a lease time when you select <b>Lease never expires</b> .

5 Click **Save changes**.

## Results

VMware Cloud Director updates the edge gateway to provide DHCP services.

## Add DHCP Bindings

If you have services running on a virtual machine and do not want the IP address to be changed, you can bind the virtual machine MAC address to the IP address. The IP address you bind must not overlap a DHCP IP pool.

## Prerequisites

You have the MAC addresses for the virtual machines for which you want to set up bindings.

## Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.

- 2 On the **DHCP > Bindings** tab, click the **Create** () button, specify the details for the binding, and click **Keep**.

Option	Description
<b>MAC Address</b>	Type the MAC address of the virtual machine that you want bound to the IP address.
<b>Host Name</b>	Type the host name you want set for that virtual machine when the virtual machine requests a DHCP lease.
<b>IP Address</b>	Type the IP address you want bound to the MAC address.
<b>Subnet Mask</b>	Type the subnet mask of the edge gateway interface.
<b>Domain Name</b>	Type the domain name of the DNS server.
<b>Auto Configure DNS</b>	Enable this toggle to use the DNS service configuration for this DNS binding. If enabled, the <b>Primary Name Server</b> and <b>Secondary Name Server</b> are set to <b>Auto</b> .
<b>Primary Name Server</b>	When you do not select <b>Auto Configure DNS</b> , type your primary DNS server IP address of your primary DNS server. This IP address is used for hostname-to-IP address resolution.
<b>Secondary Name Server</b>	When you do not select <b>Auto Configure DNS</b> , type your secondary DNS server IP address. This IP address is used for hostname-to-IP address resolution.
<b>Default Gateway</b>	Type the default gateway address. When you do not specify the default gateway IP address, the internal interface of the edge gateway instance is taken as the default gateway.
<b>Lease Never Expires</b>	Enable this toggle to keep the IP address bound to that MAC address forever. When you select this option, <b>Lease Time</b> is set to infinite.
<b>Lease Time (Seconds)</b>	Length of time (in seconds) that the DHCP-assigned IP addresses are leased to the clients. The default lease time is one day (86400 seconds).

**Note** You cannot specify a lease time when you select **Lease never expires**.

- 3 Click **Save changes**.

## Configuring DHCP Relay for NSX Data Center for vSphere Edge Gateways

The DHCP relay capability provided by NSX in your VMware Cloud Director environment enables you to leverage your existing DHCP infrastructure from within your VMware Cloud Director environment without any interruption to the IP address management in your existing DHCP infrastructure. DHCP messages are relayed from virtual machines to the designated DHCP servers in your physical DHCP infrastructure, which allows IP addresses controlled by the NSX software to continue to be synchronized with IP addresses in the rest of your DHCP-controlled environments.

The DHCP relay configuration of an edge gateway can list several DHCP servers. Requests are sent to all listed servers. While relaying the DHCP request from the VMs, the edge gateway adds a gateway IP address to the request. The external DHCP server uses this gateway address to match a pool and allocate an IP address for the request. The gateway address must belong to a subnet of the edge gateway interface.

You can specify a different DHCP server for each edge gateway and can configure multiple DHCP servers on each edge gateway to provide support for multiple IP domains.

---

**Note**

- DHCP relay does not support overlapping IP address spaces.
  - DHCP relay and DHCP service cannot run on the same vNIC at the same time. If a relay agent is configured on a vNIC, a DHCP pool cannot be configured on the subnets of that vNIC. See the *NSX Administration Guide* for details.
- 

## Specify a DHCP Relay Configuration for an NSX Data Center for vSphere Edge Gateway

The NSX software in your VMware Cloud Director environment provides the capability for the edge gateway to relay DHCP messages to DHCP servers external to your VMware Cloud Director organization virtual data center. You can configure the DHCP relay capability of the edge gateway.

As described in the *NSX Administration* documentation, the DHCP servers can be specified using an existing IP set, IP address block, domain, or a combination of all of these. DHCP messages are relayed to every specified DHCP server.

You must also configure at least one DHCP relay agent. A DHCP relay agent is an interface on the edge gateway from which the DHCP requests are relayed to the external DHCP servers.

### Prerequisites

If you want to use an IP set to specify a DHCP server, verify that an IP set exists as a grouping object available to the edge gateway. See [Create an IP Set for Use in Firewall Rules and DHCP Relay Configuration](#).

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **DHCP > Relay**.

- 3 Use the on-screen fields to specify the DHCP servers by IP addresses, domain names, or IP sets.

You select from existing IP sets using **Add** () button to browse the available IP sets.

- 4 Configure a DHCP relay agent and add its configuration to the on-screen table by clicking the **Add** () button, selecting a vNIC and its gateway IP address, and then clicking **Keep**.

By default, the Gateway IP Address matches the primary address of the selected vNIC. You can keep the default or select an alternate address if one is available on that vNIC.

- 5 Click **Save changes**.

## Managing Network Address Translation on an NSX Data Center for vSphere Edge Gateway

The NSX Data Center for vSphere software in your VMware Cloud Director environment enables the edge gateways to provide a network address translation (NAT) service. Using this capability reduces the number of public IP addresses that an organization must use, for economy and security purposes.

The edge gateway NAT service provides the ability to assign a public address to a virtual machine or group of virtual machines in a private network. To enable your edge gateways to provide access to services running on privately addressed virtual machines in your organization virtual data center, you must configure NAT rules on the edge gateways. In the most common case, you associate a NAT service with an uplink interface on an edge gateway in your VMware Cloud Director environment so that addresses on organization virtual data center networks are not exposed on the external network.

The NAT service configuration is separated into source NAT (SNAT) and destination NAT (DNAT) rules. When you configure a SNAT or a DNAT rule on an edge gateway in the VMware Cloud Director environment, you always configure the rule from the perspective of your organization virtual data center. Specifically, that means you configure the rules in the following ways:

- **SNAT:** the traffic is traveling from a virtual machine on an internal network in your organization virtual data center (the source) through the Internet to the external network (the destination). A SNAT rule translates the source IP address of the outgoing packets of an organization virtual data center network that are being sent to an external network or to another organization virtual data center network.
- **DNAT:** the traffic is traveling from the Internet (the source) to a virtual machine inside your organization virtual data center (the destination). A DNAT rule translates the IP address, and optionally the port, of packets received by an organization virtual data center network that are coming from an external network or from another organization virtual data center network.

You can configure NAT rules to create a private IP address space inside your organization virtual data center. This configuration provides the ability to port a private IP address space from one organization virtual data center to another. Configuring NAT rules allows you to use the same private IP addresses for your virtual machines in one organization virtual data center that were used in another.

The NAT rule capability in your VMware Cloud Director environment supports:

- Creating subnets within the private IP address space
- Creating multiple private IP address spaces for an edge gateway
- Configuring multiple NAT rules on multiple edge gateway interfaces

---

**Important** You must configure both firewall and NAT rules on an edge gateway for the virtual machines on an edge gateway network to be accessible. By default, edge gateways are deployed with firewall rules configured to deny all network traffic to and from the virtual machines on the edge gateway networks. Also, NAT is deactivated by default on the edge gateways so that edge gateways are unable to translate the IP addresses of the incoming and outgoing traffic unless you configure NAT on the edge gateways. Attempting to ping a virtual machine on a network after configuring a NAT rule will fail unless you add a firewall rule to allow the corresponding traffic.

---

## Add a SNAT or a DNAT Rule

You can create a source NAT (SNAT) rule to change the source IP address from a public to private IP address or the reverse. You can create a destination NAT (DNAT) rule to change the destination IP address from a public to private IP address or the reverse.

When creating NAT rules, you can specify the original and translated IP addresses by using the following formats:

- IP address; for example, 192.0.2.0
- IP address range; for example, 192.0.2.0-192.0.2.24
- IP address/subnet mask; for example, 192.0.2.0/24
- any

When you configure a SNAT or a DNAT rule on an edge gateway in the VMware Cloud Director environment, you always configure the rule from the perspective of your organization virtual data center. A SNAT rule translates the source IP address of packets sent from an organization virtual data center network out to an external network or to another organization virtual data center network. A DNAT rule translates the IP address, and optionally the port, of packets received by an organization virtual data center network that are coming from an external network or from another organization virtual data center network.

### Prerequisites

The public IP addresses must have been added to the NSX Data Center for vSphere edge gateway interface on which you want to add the rule. For DNAT rules, the original (public) IP

address must have been added to the edge gateway interface and for SNAT rules, the translated (public) IP address must have been added to the interface.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **NAT** to view the NAT Rules screen.
- 3 Depending on which type of NAT rule you are creating, click **DNAT Rule** or **SNAT Rule**.
- 4 Configure a Destination NAT rule (outside coming inside).

Option	Description
<b>Applied On</b>	Select the interface on which to apply the rule.
<b>Original IP/Range</b>	Type the required IP address or select the allocated IP address from the list. This address must be the public IP address of the edge gateway for which you are configuring the DNAT rule. In the packet being inspected, this IP address or range would be those that appear as the destination IP address of the packet. These packet destination addresses are the ones translated by this DNAT rule.
<b>Protocol</b>	Select the protocol to which the rule applies. To apply this rule on all protocols, select <b>Any</b> .
<b>Original Port</b>	(Optional) Select the port or port range that the incoming traffic uses on the edge gateway to connect to the internal network on which the virtual machines are connected. This selection is not available when the <b>Protocol</b> is set to <b>ICMP</b> or <b>Any</b> .
<b>ICMP Type</b>	When you select <b>ICMP</b> (an error reporting and a diagnostic utility used between devices to communicate error information) for <b>Protocol</b> , select the <b>ICMP Type</b> from the drop-down menu.  ICMP messages are identified by the type field. By default, the ICMP type is set to any.
<b>Translated IP/Range</b>	Type the IP address or a range of IP addresses to which destination addresses on inbound packets will be translated.  These addresses are the IP addresses of the one or more virtual machines for which you are configuring DNAT so that they can receive traffic from the external network.
<b>Translated Port</b>	(Optional) Select the port or port range that inbound traffic is connecting to on the virtual machines on the internal network. These ports are the ones into which the DNAT rule is translating for the packets inbound to the virtual machines.
<b>Source IP address</b>	If you want the rule to apply only for traffic from a specific domain, enter an IP address for this domain or an IP address range in CIDR format. If you leave this text box blank, the DNAT rule applies to all IP addresses that are in the local subnet.
<b>Source Port</b>	(Optional) Enter a port number for the source.

Option	Description
Description	(Optional) Enter a meaningful description for the DNAT rule.
Enabled	Toggle on to activate this rule.
Enable logging	Toggle on to have the address translation performed by this rule logged.

## 5 Configure a Source NAT rule (inside going outside).

Option	Description
Applied On	Select the interface on which to apply the rule.
Original Source IP/Range	Type the original IP address or range of IP addresses to apply to this rule, or select the allocated IP address from the list.  These addresses are the IP addresses of one or more virtual machines for which you are configuring the SNAT rule so that they can send traffic to the external network.
Translated Source IP/Range	Type the required IP address.  This address is always the public IP address of the gateway for which you are configuring the SNAT rule. Specifies the IP address to which source addresses (the virtual machines) on outbound packets are translated to when they send traffic to the external network.
Destination IP Address	(Optional) If you want the rule to apply only for traffic to a specific domain, enter an IP address for this domain or an IP address range in CIDR format. If you leave this text box blank, the SNAT rule applies to all destinations outside of the local subnet.
Destination Port	(Optional) Enter a port number for the destination.
Description	(Optional) Enter a meaningful description for the SNAT rule.
Enabled	Toggle on to activate this rule.
Enable logging	Toggle on to have the address translation performed by this rule logged.

6 Click **Keep** to add the rule to the on-screen table.

7 Repeat the steps to configure additional rules.

8 Click **Save changes** to save the rules to the system.

### What to do next

Add corresponding edge gateway firewall rules for the SNAT or DNAT rules you just configured. See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#).

## Advanced Routing Configuration for NSX Data Center for vSphere Edge Gateways

You can configure the static and dynamic routing on your NSX Data Center for vSphere edge gateways.

To enable dynamic routing, you configure an advanced edge gateway using the Border Gateway Protocol (BGP) or the Open Shortest Path First (OSPF) protocol.

For detailed information about the routing capabilities that NSX Data Center for vSphere provides, see the NSX Data Center for vSphere documentation.

You can specify static and dynamic routing for each advanced edge gateway. The dynamic routing capability provides the necessary forwarding information between Layer 2 broadcast domains, which allows you to decrease Layer 2 broadcast domains and improve network efficiency and scale. NSX Data Center for vSphere extends this intelligence to the locations of the workloads for East-West routing. This capability allows more direct virtual machine to virtual machine communication without the added cost or time needed to extend hops.

## Specify Default Routing Configurations for the NSX Data Center for vSphere Edge Gateway

You can specify the default settings for static routing and dynamic routing for an edge gateway.

---

**Note** To remove all configured routing settings, use the **CLEAR GLOBAL CONFIGURATION** button at the bottom of the **Routing Configuration** screen. This action deletes all routing settings currently specified on the subscreens: default routing settings, static routes, OSPF, BGP, and route redistribution.

---

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Routing > Routing Configuration**.
- 3 To enable Equal Cost Multipath (ECMP) routing for this edge gateway, turn on the **ECMP** toggle.

As described in the *NSX Administration* documentation, ECMP is a routing strategy that allows next-hop packet forwarding to a single destination to occur over multiple best paths. NSX determines these best paths either statically, using configured static routes, or as a result of metric calculations by dynamic routing protocols like OSPF or BGP. You can specify the multiple paths for static routes by specifying multiple next hops on the Static Routes screen.

For more details about ECMP and NSX, see the routing topics in the *NSX Troubleshooting Guide*.

- 4 Specify settings for the default routing gateway.
  - a Use the **Applied On** drop-down list to select an interface from which the next hop towards the destination network can be reached.
 

To see details about the selected interface, click the blue information icon.
  - b Type the gateway IP address.
  - c Type the MTU.

- d (Optional) Type an optional description.
- e Click **Save changes**.

## 5 Specify default dynamic routing settings.

---

**Note** If you have IPsec VPN configured in your environment, you should not use dynamic routing.

---

- a Select a router ID.

You can select a router ID in the list or use the + icon to enter a new one. This router ID is the first uplink IP address of the edge gateway that pushes routes to the kernel for dynamic routing.

- b Configure logging by turning on the **Enable Logging** toggle and selecting the log level.
- c Click **OK**.

## 6 Click **Save changes**.

### What to do next

Add static routes. See [Add a Static Route](#).

Configure route redistribution. See [Configure Route Redistributions](#).

Configure dynamic routing. See the following topics:

- [Configure BGP](#)
- [Configure OSPF](#)

## Add a Static Route

You can add a static route for a destination subnet or host.

If ECMP is enabled in the default routing configuration, you can specify multiple next hops in the static routes. See [Specify Default Routing Configurations for the NSX Data Center for vSphere Edge Gateway](#) for steps on enabling ECMP.

### Prerequisites

As described in the NSX documentation, the next hop IP address of the static route must exist in a subnet associated with one of the NSX Data Center for vSphere edge gateway interfaces. Otherwise, configuration of that static route fails.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Routing > Static Routes**.

- 3 Click the **Create** () button.
- 4 Configure the following options for the static route:

Option	Description
<b>Network</b>	Type the network in CIDR notation.
<b>Next Hop</b>	Type the IP address of the next hop. The next hop IP address must exist in a subnet associated with one of the edge gateway interfaces. If ECMP is enabled, you can type multiple next hops.
<b>MTU</b>	Edit the maximum transmission value for data packets. The MTU value cannot be higher than the MTU value set on the selected edge gateway interface. You can see the MTU set on the edge gateway interface by default on the Routing Configuration screen.
<b>Interface</b>	Optionally, select the edge gateway interface on which you want to add a static route. By default, the interface is selected that matches the next hop address.
<b>Description</b>	Optionally, type a description for the static route.

- 5 Click **Save changes**.

#### What to do next

Configure a NAT rule for the static route. See [Add a SNAT or a DNAT Rule](#).

Add a firewall rule to allow traffic to traverse the static route. See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#).

## Configure OSPF

You can configure the Open Shortest Path First (OSPF) routing protocol for the dynamic routing capabilities of an NSX Data Center for vSphere edge gateway. A common application of OSPF on an edge gateway in a VMware Cloud Director environment is to exchange routing information between edge gateways in VMware Cloud Director.

The NSX edge gateway supports OSPF, an interior gateway protocol that routes IP packets only within a single routing domain. As described in the *NSX Administration* documentation, configuring OSPF on an NSX edge gateway enables the edge gateway to learn and advertise routes. The edge gateway uses OSPF to gather link state information from available edge gateways and construct a topology map of the network. The topology determines the routing table presented to the Internet layer, which makes routing decisions based on the destination IP address found in IP packets.

As a result, OSPF routing policies provide a dynamic process of traffic load balancing between routes of equal cost. An OSPF network is divided into routing areas to optimize traffic flow and limit the size of routing tables. An area is a logical collection of OSPF networks, routers, and links that have the same area identification. Areas are identified by an Area ID.

## Prerequisites

A Router ID must be configured . [Specify Default Routing Configurations for the NSX Data Center for vSphere Edge Gateway.](#)

## Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Routing > OSPF**.
- 3 If OSPF is not currently enabled, use the **OSPF Enabled** toggle to enable it.
- 4 Configure the OSPF settings according to the needs of your organization.

Option	Description
<b>Enable Graceful Restart</b>	Specifies that packet forwarding is to remain uninterrupted when OSPF services are restarted.
<b>Enable Default Originate</b>	Allows the edge gateway to advertise itself as a default gateway to its OSPF peers.

- 5 (Optional) You can either click **Save changes** or continue with configuring area definitions and interface mappings.

- 6 Add an OSPF area definition by clicking the **Add** () button, specifying details for the mapping in the dialog box, and clicking **Keep**.

**Note** By default, the system configures a not-so-stubby area (NSSA) with area ID of 51, and this area is automatically displayed in the area definitions table on the OSPF screen. You can modify or delete the NSSA area.

Option	Description
<b>Area ID</b>	Type an area ID in the form of an IP address or decimal number.
<b>Area Type</b>	Select <b>Normal</b> or <b>NSSA</b> . NSSAs prevent the flooding of AS-external link-state advertisements (LSAs) into NSSAs. They rely on default routing to external destinations. As a result, NSSAs must be placed at the edge of an OSPF routing domain. NSSA can import external routes into the OSPF routing domain, by that means providing transit service to small routing domains that are not part of the OSPF routing domain.
<b>Area Authentication</b>	Select the type of authentication for OSPF to perform at the area level. All edge gateways within the area must have the same authentication and corresponding password configured. For MD5 authentication to work, both the receiver and transmitter must have the same MD5 key. Choices are: <ul style="list-style-type: none"> <li>■ <b>None</b> No authentication is required.</li> <li>■ <b>Password</b> With this choice, the password you specify in the <b>Area Authentication Value</b> field is included in the transmitted packet.</li> <li>■ <b>MD5</b> With this choice, the authentication uses MD5 (Message Digest type 5) encryption. An MD5 checksum is included in the transmitted packet. Type the Md5 key into the <b>Area Authentication Value</b> field.</li> </ul>

- 7 Click **Save changes**, so that the newly configured area definitions are available for selection when you add interface mappings.
- 8 Add an interface mapping by clicking the **Add** () button, specifying details for the mapping in the dialog box, and clicking **Keep**.

These mappings map the edge gateway interfaces to the areas.

- a In the dialog box, select the interface you want to map to an area definition.

The interface specifies the external network that both edge gateways are connected to.

- b Select the area ID for the area to map to the selected interface.

- c (Optional) Change the OSPF settings from the default values to customize them for this interface mapping.

When configuring a new mapping, the default values for these settings are displayed. In most cases, it is recommended to retain the default settings. If you do change the settings, make sure that the OSPF peers use the same settings.

Option	Description
Hello Interval	Interval (in seconds) between hello packets that are sent on the interface.
Dead Interval	Interval (in seconds) during which at least one hello packet must be received from a neighbor before that neighbor is declared down.
Priority	Priority of the interface. The interface with the highest priority is the designated edge gateway router.
Cost	Overhead required to send packets across that interface. The cost of an interface is inversely proportional to the bandwidth of that interface. The larger the bandwidth, the smaller the cost.

- d Click **Keep**.

- 9 Click **Save changes** in the OSPF screen.

#### What to do next

Configure OSPF on the other edge gateways that you want to exchange routing information with.

Add a firewall rule that allows traffic between the OSPF-enabled edge gateways. See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#).

Make sure that the route redistribution and firewall configuration allow the correct routes to be advertised. See [Configure Route Redistributions](#).

## Configure BGP

You can configure Border Gateway Protocol (BGP) for the dynamic routing capabilities of an NSX Data Center for vSphere edge gateway.

As described in the *NSX Administration Guide*, BGP makes core routing decisions by using a table of IP networks or prefixes, which designate network reachability among multiple autonomous systems. In the networking field, the term BGP speaker refers to a networking device that is running BGP. Two BGP speakers establish a connection before any routing information is exchanged. The term BGP neighbor refers to a BGP speaker that has established such a connection. After establishing the connection, the devices exchange routes and synchronize their tables. Each device sends keep alive messages to keep this relationship alive.

## Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Routing > BGP**.
- 3 If BGP is not currently enabled, use the **Enable BGP** toggle to enable it.
- 4 Configure the BGP settings according to the needs of your organization.

Option	Description
<b>Enable Graceful Restart</b>	Specifies that packet forwarding is to remain uninterrupted when BGP services are restarted.
<b>Enable Default Originate</b>	Allows the edge gateway to advertise itself as a default gateway to its BGP neighbors.
<b>Local AS</b>	<p>Required. Specify the autonomous system (AS) ID number to use for the local AS feature of the protocol. The value you specify must be a globally unique number between 1 and 65534.</p> <p>The local AS is a feature of BGP. The system assigns the local AS number to the edge gateway you are configuring. The edge gateway advertises this ID when the edge gateway peers with its BGP neighbors in other autonomous systems. The path of autonomous systems that a route would traverse is used as one metric in the dynamic routing algorithm when selecting the best path to a destination.</p>

- 5 You can either click **Save changes**, or continue to configure settings for the BGP routing neighbors.
- 6 Add a BGP neighbor configuration by clicking the **Add** () button, specifying details for the neighbor in the dialog box, and clicking **Keep**.

Option	Description
<b>IP Address</b>	Type the IP address of a BGP neighbor for this edge gateway.
<b>Remote AS</b>	Type a globally unique number between 1-65534 for the autonomous system to which this BGP neighbor belongs. This remote AS number is used in the BGP neighbor's entry in the system's BGP neighbors table.
<b>Weight</b>	The default weight for the neighbor connection. Adjust as appropriate for your organization's needs.
<b>Keep Alive Time</b>	The frequency with which the software sends keep alive messages to its peer. The default frequency is 60 seconds. Adjust as appropriate for the needs of your organization.

Option	Description
<b>Hold Down Time</b>	<p>The interval for which the software declares a peer dead after not receiving a keep alive message. This interval must be three times the keep alive interval. The default interval is 180 seconds. Adjust as appropriate for the needs of your organization.</p> <p>Once peering between two BGP neighbors is achieved, the edge gateway starts a hold down timer. Every keep alive message it receives from the neighbor resets the hold down timer to 0. If the edge gateway fails to receive three consecutive keep alive messages, so that the hold down timer reaches three times the keep alive interval, the edge gateway considers the neighbor down and deletes the routes from this neighbor.</p>
<b>Password</b>	<p>If this BGP neighbor requires authentication, type the authentication password.</p> <p>Each segment sent on the connection between the neighbors is verified. MD5 authentication must be configured with the same password on both BGP neighbors, otherwise, the connection between them will not be made.</p>
<b>BGP Filters</b>	<p>Use this table to specify route filtering using a prefix list from this BGP neighbor.</p> <p><b>Caution</b> A <code>block all</code> rule is enforced at the end of the filters.</p> <p>Add a filter to the table by clicking the + icon and configuring the options. Click <b>Keep</b> to save each filter.</p> <ul style="list-style-type: none"> <li>■ Select the direction to indicate whether you are filtering traffic to or from the neighbor.</li> <li>■ Select the action to indicate whether you are allowing or denying traffic.</li> <li>■ Type the network that you want to filter to or from the neighbor. Type <code>ANY</code> or a network in a CIDR format.</li> <li>■ Type the <b>IP Prefix GE</b> and <b>IP Prefix LE</b> to use the <code>le</code> and <code>ge</code> keywords in the IP prefix list.</li> </ul>

7 Click **Save changes** to save the configurations to the system.

#### What to do next

Configure BGP on the other edge gateways that you want to exchange routing information with.

Add a firewall rule that allows traffic to and from the BGP-configured edge gateways. See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#) for information.

## Configure Route Redistributions

By default the router only shares routes with other routers running the same protocol. When you have configured a multi-protocol environment, you must configure route redistribution to have cross-protocol route sharing. You can configure route redistribution for an NSX Data Center for vSphere edge gateway.

**Procedure**

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Routing > Route Redistribution**.
- 3 Use the protocol toggles to turn on those protocols for which you want to enable route redistribution.
- 4 Add IP prefixes to the on-screen table.

- a Click the **Add** () button.
- b Type a name and the IP address of the network in CIDR format.
- c Click **Keep**.

- 5 Specify redistribution criteria for each IP prefix by clicking the **Add** () button, specifying the criteria in the dialog box, and clicking **Keep**.

Entries in the table are processed sequentially. Use the up and down arrows to adjust the sequence.

Option	Description
<b>Prefix Name</b>	Select a specific IP prefix to apply this criteria to or select <b>Any</b> to apply the criteria to all network routes.
<b>Learner Protocol</b>	Select the protocol that is to learn routes from other protocols under this redistribution criteria.
<b>Allow learning from</b>	Select the types of networks from which routes can be learned for the protocol selected in the <b>Learner Protocol</b> list.
<b>Action</b>	Select whether to permit or deny redistribution from the selected types of networks.

- 6 Click **Save changes**.

## Load Balancing with NSX Data Center for vSphere

The load balancer distributes incoming service requests among multiple servers in such a way that the load distribution is transparent to users. Load balancing provides application high availability and helps achieve optimal resource utilization, maximizing throughput, minimizing response time, and avoiding overload.

### About Load Balancing

The load balancer distributes incoming service requests among multiple servers in such a way that the load distribution is transparent to users. Load balancing helps achieve optimal resource use, maximizing throughput, minimizing response time, and avoiding overload.

The NSX load balancer supports two load balancing engines. The layer 4 load balancer is packet-based and provides fast-path processing. The layer 7 load balancer is socket-based and supports advanced traffic management strategies and DDOS mitigation for back end services.

Load balancing for an NSX Data Center for vSphere edge gateway is configured on the external interface because the edge gateway load balances incoming traffic from the external network. When configuring virtual servers for load balancing, specify one of the available IP addresses you have in your organization VDC.

### **Load Balancing Strategies and Concepts**

A packet-based load balancing strategy is implemented on the TCP and UDP layer. Packet-based load balancing does not stop the connection or buffer the whole request. Instead, after manipulating the packet, it sends it directly to the selected server. TCP and UDP sessions are maintained in the load balancer so that packets for a single session are directed to the same server. You can select Acceleration Enable in both the global configuration and relevant virtual server configuration to enable packet-based load balancing.

A socket-based load balancing strategy is implemented on top of the socket interface. Two connections are established for a single request, a client-facing connection and a server-facing connection. The server-facing connection is established after server selection. For the HTTP socket-based implementation, the whole request is received before sending to the selected server with optional L7 manipulation. For HTTPS socket-based implementation, authentication information is exchanged either on the client-facing connection or server-facing connection. Socket-based load balancing is the default mode for TCP, HTTP, and HTTPS virtual servers.

The key concepts of the NSX load balancer are, virtual server, server pool, server pool member, and service monitor.

#### **Virtual Server**

Abstract of an application service, represented by a unique combination of IP, port, protocol and application profile such as TCP or UDP.

#### **Server Pool**

Group of back end servers.

#### **Server Pool Member**

Represents the back end server as member in a pool.

#### **Service Monitor**

Defines how to probe the health status of a back end server.

#### **Application Profile**

Represents the TCP, UDP, persistence, and certificate configuration for a given application.

## Setup Overview

You begin by setting global options for the load balancer. You now create a server pool consisting of back end server members and associate a service monitor with the pool to manage and share the back end servers efficiently.

You then create an application profile to define the common application behavior in a load balancer such as client SSL, server SSL, x-forwarded-for, or persistence. Persistence sends subsequent requests with similar characteristic such as, source IP or cookie are required to be dispatched to the same pool member, without running the load balancing algorithm. The application profile can be reused across virtual servers.

You then create an optional application rule to configure application-specific settings for traffic manipulation such as, matching a certain URL or hostname so that different requests can be handled by different pools. Next, you create a service monitor that is specific to your application or you can use an existing service monitor if it meets your needs.

Optionally you can create an application rule to support advanced functionality of L7 virtual servers. Some use cases for application rules include content switching, header manipulation, security rules, and DOS protection.

Finally, you create a virtual server that connects your server pool, application profile, and any potential application rules together.

When the virtual server receives a request, the load balancing algorithm considers pool member configuration and runtime status. The algorithm then calculates the appropriate pool to distribute the traffic comprising one or more members. The pool member configuration includes settings such as, weight, maximum connection, and condition status. The runtime status includes current connections, response time, and health check status information. The calculation methods can be round-robin, weighted round-robin, least connection, source IP hash, weighted least connections, URL, URI, or HTTP header.

Each pool is monitored by the associated service monitor. When the load balancer detects a problem with a pool member, it is marked as DOWN. Only UP server is selected when choosing a pool member from the server pool. If the server pool is not configured with a service monitor, all the pool members are considered as UP.

## Configure the Load Balancer Service

Global load balancer configuration parameters include overall enablement, selection of the layer 4 or layer 7 engine, and specification of the types of events to log.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Load Balancer > Global Configuration**.

### 3 Select the options you want to enable:

Option	Action
Status	<p>Enable the load balancer by clicking the toggle icon.</p> <p>Enable <b>Acceleration Enabled</b> to configure the load balancer to use the faster L4 engine rather than L7 engine. The L4 TCP VIP is processed before the edge gateway firewall so no Allow firewall rule is required.</p> <hr/> <p><b>Note</b> L7 VIPs for HTTP and HTTPS are processed after the firewall, so when you do not enable acceleration, an edge gateway firewall rule must exist to allow access to the L7 VIP for those protocols. When you enable acceleration, and the server pool is in a non-transparent mode, a SNAT rule is added, so you must ensure that the firewall is enabled on the edge gateway.</p> <hr/>
Enable Logging	Enable logging so that the edge gateway load balancer collects traffic logs.
Log Level	Choose the severity of events to be collected in the logs.

### 4 Click **Save changes**.

#### What to do next

Configure application profiles for the load balancer. See [Create an Application Profile](#).

#### Create an Application Profile

An application profile defines the behavior of the load balancer for a particular type of network traffic. After configuring a profile, you associate it with a virtual server. The virtual server then processes traffic according to the values specified in the profile. Using profiles enhances your control over managing network traffic, and makes traffic-management tasks easier and more efficient.

When you create a profile for HTTPS traffic, the following HTTPS traffic patterns are allowed:

- Client -> HTTPS -> LB (terminate SSL) -> HTTP -> servers
- Client -> HTTPS -> LB (terminate SSL) -> HTTPS -> servers
- Client -> HTTPS-> LB (SSL passthrough) -> HTTPS -> servers
- Client -> HTTP-> LB -> HTTP -> servers

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Load Balancer > Application Profiles**.
- 3 Click the **Create** () button.
- 4 Enter a name for the profile.

## 5 Configure the application profile.

Option	Description
<b>Type</b>	Select the protocol type used to send requests to the server. The list of required parameters depends on the protocol you select. Parameters that are not applicable to the protocol you selected cannot be entered. All other parameters are required.
<b>Enable SSL Passthrough</b>	Click to enable SSL authentication to be passed through to the virtual server. Otherwise SSL authentication takes place at the destination address.
<b>HTTP Redirect URL</b>	(HTTP and HTTPS) Enter the URL to which traffic that arrives at the destination address should be redirected.
<b>Persistence</b>	<p>Specify a persistence mechanism for the profile.</p> <p>Persistence tracks and stores session data, such as the specific pool member that serviced a client request. This ensures that client requests are directed to the same pool member throughout the life of a session or during subsequent sessions. The options are:</p> <ul style="list-style-type: none"> <li>■ <b>Source IP</b> <p>Source IP persistence tracks sessions based on the source IP address. When a client requests a connection to a virtual server that supports source address affinity persistence, the load balancer checks to see if that client previously connected, and if so, returns the client to the same pool member.</p> </li> <li>■ <b>MSRDP</b> <p>(TCP Only) Microsoft Remote Desktop Protocol persistence (MSRDP) maintains persistent sessions between Windows clients and servers that are running the Microsoft Remote Desktop Protocol (RDP) service. The recommended scenario for enabling MSRDP persistence is to create a load balancing pool that consists of members running a Windows Server guest OS, where all members belong to a Windows cluster and participate in a Windows session directory.</p> </li> <li>■ <b>SSL Session ID</b> <p>SSL Session ID persistence is available when you enable SSL passthrough. SSL Session ID persistence ensures that repeat connections from the same client are sent to the same server. Session ID persistence allows the use of SSL session resumption, which saves processing time for both the client and the server.</p> </li> </ul>
<b>Cookie Name</b>	(HTTP and HTTPS) If you specified <b>Cookie</b> as the persistence mechanism, enter the cookie name. Cookie persistence uses a cookie to uniquely identify the session the first time a client accesses the site. The load balancer refers to this cookie when connecting subsequent requests in the session, so that they all go to the same virtual server.

Option	Description
Mode	<p>Select the mode by which the cookie should be inserted. The following modes are supported:</p> <ul style="list-style-type: none"> <li> <b>Insert</b> <p>The edge gateway sends a cookie. When the server sends one or more cookies, the client will receive one extra cookie (the server cookies plus the edge gateway cookie). When the server does not send any cookies, the client will receive the edge gateway cookie only.</p> </li> <li> <b>Prefix</b> <p>Select this option when your client does not support more than one cookie.</p> <p><b>Note</b> All browsers accept multiple cookies. But you might have a proprietary application using a proprietary client that supports only one cookie. The Web server sends its cookie as usual. The edge gateway injects (as a prefix) its cookie information in the server cookie value. This cookie added information is removed when the edge gateway sends it to the server.</p> </li> <li> <b>App Session</b> For this option, the server does not send a cookie. Instead, it sends the user session information as a URL. For example, <code>http://example.com/admin/UpdateUserServlet;jsessionid=0I24B9ASD7BSSD</code>, where <code>jsessionid</code> is the user session information and is used for the persistence. It is not possible to see the App Session persistence table for troubleshooting. </li> </ul>
Expires in (Seconds)	<p>Enter a length of time in seconds that persistence stays in effect. Must be a positive integer in the range 1–86400.</p> <p><b>Note</b> For L7 load balancing using TCP source IP persistence, the persistence entry times out if no new TCP connections are made for a period of time, even if the existing connections are still alive.</p>
Insert X-Forwarded-For HTTP header	<p>(HTTP and HTTPS) Select <b>Insert X-Forwarded-For HTTP</b> header for identifying the originating IP address of a client connecting to a Web server through the load balancer.</p> <p><b>Note</b> Using this header is not supported if you enabled SSL passthrough.</p>
Enable Pool Side SSL	<p>(HTTPS Only) Select <b>Enable Pool Side SSL</b> to define the certificate, CAs, or CRLs used to authenticate the load balancer from the server side in the Pool Certificates tab.</p>

- 6 (HTTPS only) Configure the certificates to be used with the application profile. If the certificates you need do not exist, you can create them from the **Certificates** tab.

Option	Description
Virtual Server Certificates	Select the certificate, CAs, or CRLs used to decrypt HTTPS traffic.
Pool Certificates	<p>Define the certificate, CAs, or CRLs used to authenticate the load balancer from the server side.</p> <p><b>Note</b> Select <b>Enable Pool Side SSL</b> to enable this tab.</p>

Option	Description
Cipher	Select the cipher algorithms (or cipher suite) negotiated during the SSL/TLS handshake.
Client Authentication	Specify whether client authentication is to be ignored or required. <b>Note</b> When set to <b>Required</b> , the client must provide a certificate after the request or the handshake is canceled.

7 To preserve your changes, click **Keep**.

### What to do next

Add service monitors for the load balancer to define health checks for different types of network traffic. See [Create a Service Monitor](#).

### Create a Service Monitor

You create a service monitor to define health check parameters for a particular type of network traffic. When you associate a service monitor with a pool, the pool members are monitored according to the service monitor parameters.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Load Balancer > Service Monitoring**.
- 3 Click the **Create** () button.
- 4 Enter a name for the service monitor.
- 5 (Optional) Configure the following options for the service monitor:

Option	Description
Interval	Enter the interval at which a server is to be monitored using the specified <b>Method</b> .
Timeout	Enter the maximum time in seconds within which a response from the server must be received.
Max Retries	Enter the number of times the specified monitoring <b>Method</b> must fail sequentially before the server is declared down.
Type	Select the way in which you want to send the health check request to the server—HTTP, HTTPS, TCP, ICMP, or UDP. Depending on the type selected, the remaining options in the <b>New Service Monitor</b> dialog are activated or deactivated.
Expected	(HTTP and HTTPS) Enter the string that the monitor expects to match in the status line of the HTTP or HTTPS response (for example, HTTP/1.1).

Option	Description
Method	(HTTP and HTTPS) Select the method to be used to detect server status.
URL	(HTTP and HTTPS) Enter the URL to be used in the server status request. <b>Note</b> When you select the POST method, you must specify a value for <b>Send</b> .
Send	(HTTP, HTTPS, UDP) Enter the data to be sent.
Receive	(HTTP, HTTPS, and UDP) Enter the string to be matched in the response content. <b>Note</b> When <b>Expected</b> is not matched, the monitor does not try to match the <b>Receive</b> content.
Extension	(ALL) Enter advanced monitor parameters as key=value pairs. For example, warning=10 indicates that when a server does not respond within 10 seconds, its status is set as warning. All extension items should be separated with a carriage return character. For example: <pre>&lt;extension&gt;delay=2 critical=3 escape&lt;/extension&gt;</pre>

6 To preserve your changes, click **Keep**.

#### Example: Extensions Supported for Each Protocol

Table 5-4. Extensions for HTTP/HTTPS Protocols

Monitor Extension	Description
no-body	Does not wait for a document body and stops reading after the HTTP/HTTPS header. <b>Note</b> An HTTP GET or HTTP POST is still sent; not a HEAD method.
max-age= <i>SECONDS</i>	Warns when a document is more than <i>SECONDS</i> old. The number can be in the form 10m for minutes, 10h for hours, or 10d for days.
content-type= <i>STRING</i>	Specifies a Content-Type header media type in POST calls.
linespan	Allows regex to span newlines (must precede -r or -R).
regex= <i>STRING</i> or ereg= <i>STRING</i>	Searches the page for regex <i>STRING</i> .
eregi= <i>STRING</i>	Searches the page for case-insensitive regex <i>STRING</i> .
invert-regex	Returns CRITICAL when found and OK when not found.
proxy-authorization= <i>AUTH_PAIR</i>	Specifies the username:password on proxy servers with basic authentication.
useragent= <i>STRING</i>	Sends the string in the HTTP header as User Agent.

Table 5-4. Extensions for HTTP/HTTPS Protocols (continued)

Monitor Extension	Description
header= <i>STRING</i>	Sends any other tags in the HTTP header. Use multiple times for additional headers.
onredirect=ok warning critical follow sticky stickyport	Indicates how to handle redirected pages. <i>sticky</i> is like <i>follow</i> but stick to the specified IP address. <i>stickyport</i> ensures the port stays the same.
pagesize= <i>INTEGER:INTEGER</i>	Specifies the minimum and maximum page sizes required in bytes.
warning= <i>DOUBLE</i>	Specifies the response time in seconds to result in a warning status.
critical= <i>DOUBLE</i>	Specifies the response time in seconds to result in a critical status.

Table 5-5. Extensions for HTTPS Protocol Only

Monitor Extension	Description
sni	Enables SSL/TLS hostname extension support (SNI).
certificate= <i>INTEGER</i>	Specifies the minimum number of days a certificate has to be valid. The port defaults to 443. When this option is used, the URL is not checked.
authorization=AUTH_PAIR	Specifies the username:password on sites with basic authentication.

Table 5-6. Extensions for TCP Protocol

Monitor Extension	Description
escape	Allows for the use of <code>\n</code> , <code>\r</code> , <code>\t</code> , or <code>\</code> in a send or quit string. Must come before a send or quit option. By default, nothing is added to send and <code>\r\n</code> is added to the end of quit.
all	Specifies all expect strings need to occur in a server response. By default, <code>any</code> is used.
quit= <i>STRING</i>	Sends a string to the server to cleanly close the connection.
refuse=ok warn crit	Accepts TCP refusals with states <code>ok</code> , <code>warn</code> , or <code>crit</code> . By default, uses state <code>crit</code> .
mismatch=ok warn crit	Accepts expected string mismatches with states <code>ok</code> , <code>warn</code> , or <code>crit</code> . By default, uses state <code>warn</code> .
jail	Hides output from the TCP socket.
maxbytes= <i>INTEGER</i>	Closes the connection when more than the specified number of bytes are received.
delay= <i>INTEGER</i>	Waits the specified number of seconds between sending the string and polling for a response.

Table 5-6. Extensions for TCP Protocol (continued)

Monitor Extension	Description
certificate= <i>INTEGER</i> [, <i>INTEGER</i> ]	Specifies the minimum number of days a certificate has to be valid. The first value is #days for warning and the second value is critical (if not specified - 0).
ssl	Uses SSL for the connection.
warning= <i>DOUBLE</i>	Specifies the response time in seconds to result in a warning status.
critical= <i>DOUBLE</i>	Specifies the response time in seconds to result in a critical status.

### What to do next

Add server pools for your load balancer. See [Add a Server Pool for Load Balancing](#).

### Add a Server Pool for Load Balancing

You can add a server pool to manage and share backend servers flexibly and efficiently. A pool manages load balancer distribution methods and has a service monitor attached to it for health check parameters.

#### Procedure

- Open Edge Gateway Services.
  - In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - Select the edge gateway that you want to edit and click **Services**.
- Navigate to **Load Balancer > Pools**.
- Click the **Create** () button.
- Type a name and, optionally, a description for the load balancer pool.
- Select a balancing method for the service from the **Algorithm** drop-down menu:

Option	Description
ROUND-ROBIN	Each server is used in turn according to the weight assigned to it. This is the smoothest and fairest algorithm when the server processing time remains equally distributed.
IP-HASH	Selects a server based on a hash of the source and destination IP address of each packet.
LEASTCONN	Distributes client requests to multiple servers based on the number of connections already open on the server. New connections are sent to the server with the fewest open connections.

Option	Description
URI	The left part of the URI (before the question mark) is hashed and divided by the total weight of the running servers. The result designates which server will receive the request. This option ensures that a URI is always directed to the same server as long as the server does not go down.
HTTPHEADER	HTTP header name is looked up in each HTTP request. The header name in parenthesis is not case sensitive which is similar to the ACL 'hdr()' function. If the header is absent or does not contain any value, the round robin algorithm is applied. The HTTP HEADER algorithm parameter has one option <code>headerName=&lt;name&gt;</code> . For example, you can use <code>host</code> as the HTTP HEADER algorithm parameter.
URL	URL parameter specified in the argument is looked up in the query string of each HTTP GET request. If the parameter is followed by an equal sign = and a value, then the value is hashed and divided by the total weight of the running servers. The result designates which server receives the request. This process is used to track user identifiers in requests and ensure that a same user ID is always sent to the same server as long as no server goes up or down. If no value or parameter is found, then a round robin algorithm is applied. The URL algorithm parameter has one option <code>urlParam=&lt;url&gt;</code> .

## 6 Add members to the pool.

- a Click the **Add** () button.
- b Enter the name for the pool member.
- c Enter the IP address of the pool member.
- d Enter the port at which the member is to receive traffic from the load balancer.
- e Enter the monitor port at which the member is to receive health monitor requests.
- f In the **Weight** text box, type the proportion of traffic this member is to handle. Must be an integer in the range 1-256.
- g (Optional) In the **Max Connections** text box, type the maximum number of concurrent connections the member can handle.

When the number of incoming requests exceeds the maximum, requests are queued and the load balancer waits for a connection to be released.

- h (Optional) In the **Min Connections** text box, type the minimum number of concurrent connections a member must always accept.
- i Click **Keep** to add the new member to the pool.

The operation can take a minute to complete.

## 7 (Optional) To make client IP addresses visible to the back end servers, select **Transparent**.

When **Transparent** is not selected (the default value), back end servers see the IP address of the traffic source as the internal IP address of the load balancer.

When **Transparent** is selected, the source IP address is the actual IP address of the client and the edge gateway must be set as the default gateway to ensure that return packets go through the edge gateway.

- 8 To preserve your changes, click **Keep**.

#### What to do next

Add virtual servers for your load balancer. A virtual server has a public IP address and services all incoming client requests. See [Add a Virtual Server](#).

#### Add an Application Rule

You can write an application rule to directly manipulate and manage IP application traffic.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Load Balancer > Application Rules**.

- 3 Click the **Add** () button.

- 4 Enter the name for the application rule.
- 5 Enter the script for the application rule.

For information on the application rule syntax, see <http://cbonte.github.io/haproxy-dconv/2.2/configuration.html>.

- 6 To preserve your changes, click **Keep**.

#### What to do next

Associate the new application rule to a virtual server added for the load balancer. See [Add a Virtual Server](#).

#### Add a Virtual Server

Add an NSX Data Center for vSphere edge gateway internal or uplink interface as a virtual server. A virtual server has a public IP address and services all incoming client requests.

By default, the load balancer closes the server TCP connection after each client request.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **Load Balancer > Virtual Servers**.

3 Click the **Add** () button.

4 On the **General** tab, configure the following options for the virtual server:

Option	Description
<b>Enable Virtual Server</b>	Click to enable the virtual server.
<b>Enable Acceleration</b>	Click to enable acceleration.
<b>Application Profile</b>	Select an application profile to be associated with the virtual server.
<b>Name</b>	Type a name for the virtual server.
<b>Description</b>	Type an optional description for the virtual server.
<b>IP Address</b>	Type or browse to select the IP address that the load balancer listens on.
<b>Protocol</b>	Select the protocol that the virtual server accepts. You must select the same protocol used by the selected <b>Application Profile</b> .
<b>Port</b>	Type the port number that the load balancer listens on.
<b>Default Pool</b>	Choose the server pool that the load balancer will use.
<b>Connection Limit</b>	(Optional) Type the maximum concurrent connections that the virtual server can process.
<b>Connection Rate Limit (CPS)</b>	(Optional) Type the maximum incoming new connection requests per second.

5 (Optional) To associate application rules with the virtual server, click the **Advanced** tab and complete the following steps:

a Click the **Add** () button.

The application rules created for the load balancer appear. If necessary, add application rules for the load balancer. See [Add an Application Rule](#).

6 To preserve your changes, click **Keep**.

#### What to do next

Create an edge gateway firewall rule to permit traffic to the new virtual server (the destination IP address). See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#)

## Configure Secure Access Using VPN on an NSX Data Center for vSphere Edge Gateway

You can configure the VPN capabilities that are provided by the NSX Data Center for vSphere software on your NSX Data Center for vSphere edge gateways. You can configure VPN connections to your organization virtual data center using an SSL VPN-Plus tunnel, an IPsec VPN tunnel, or an L2 VPN tunnel.

As described in the *NSX Administration Guide*, the NSX edge gateway supports these VPN services:

- SSL VPN-Plus, which allows remote users to access private corporate applications.
- IPsec VPN, which offers site-to-site connectivity between an NSX edge gateway and remote sites which also have NSX or which have third-party hardware routers or VPN gateways.
- L2 VPN, which allows extension of your organization virtual data center by allowing virtual machines to retain network connectivity while retaining the same IP address across geographical boundaries.

In a VMware Cloud Director environment, you can create VPN tunnels between:

- Organization virtual data center networks on the same organization
- Organization virtual data center networks on different organizations
- Between an organization virtual data center network and an external network

---

**Note** VMware Cloud Director does not support multiple VPN tunnels between the same two edge gateways. If there is an existing tunnel between two edge gateways and you want to add another subnet to the tunnel, delete the existing VPN tunnel and create a new one that includes the new subnet.

---

After you configure VPN tunnels for an edge gateway, you can use a VPN client from a remote location to connect to the organization virtual data center that is backed by that edge gateway.

## Configure SSL VPN-Plus

The SSL VPN-Plus services for an NSX Data Center for vSphere edge gateway in a VMware Cloud Director environment enable remote users to connect securely to the private networks and applications in the organization virtual data centers backed by that edge gateway. You can configure various SSL VPN-Plus services on the edge gateway.

In your VMware Cloud Director environment, the edge gateway SSL VPN-Plus capability supports network access mode. Remote users must install an SSL client to make secure connections and access the networks and applications behind the edge gateway. As part of the edge gateway SSL VPN-Plus configuration, you add the installation packages for the operating system and configure certain parameters. See [Add an SSL VPN-Plus Client Installation Package](#) for details.

Configuring SSL VPN-Plus on an edge gateway is a multi-step process.

## Prerequisites

Verify that all SSL certificates needed for the SSL VPN-Plus have been added to the **Certificates** screen. See [SSL Certificate Management on an NSX Data Center for vSphere Edge Gateway](#).

---

**Note** On an edge gateway, port 443 is the default port for HTTPS. For the SSL VPN functionality, the edge gateway HTTPS port must be accessible from external networks. The SSL VPN client requires the edge gateway IP address and port that are configured in the Server Settings screen on the **SSL VPN-Plus** tab to be reachable from the client system. See [Configure SSL VPN Server Settings](#).

---

## Procedure

### 1 [Navigate to the SSL-VPN Plus Screen](#)

You can navigate to the SSL-VPN Plus screen to begin configuring the SSL-VPN Plus service for an NSX Data Center for vSphere edge gateway.

### 2 [Configure SSL VPN Server Settings](#)

These server settings configure the SSL VPN server, such as the IP address and port the service listens on, the cipher list of the service, and its service certificate. When connecting to the NSX Data Center for vSphere edge gateway, remote users specify the same IP address and port you set in these server settings.

### 3 [Create an IP Pool for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#)

The remote users are assigned virtual IP addresses from the static IP pools that you configure using the **IP Pools** screen on the **SSL VPN-Plus** tab.

### 4 [Add a Private Network for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#)

Use the Private Networks screen on the **SSL VPN-Plus** tab to configure the private networks. The private networks are the ones you want the VPN clients to have access to, when the remote users connect using their VPN clients and the SSL VPN tunnel. The activated private networks will be installed in the routing table of the VPN client.

### 5 [Configure an Authentication Service for SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#)

Use the **Authentication** screen on the **SSL VPN-Plus** tab to set up a local authentication server for the edge gateway SSL VPN service and optionally enable client certificate authentication. This authentication server is used to authenticate the connecting users. All users configured in the local authentication server will be authenticated.

### 6 [Add SSL VPN-Plus Users to the Local SSL VPN-Plus Authentication Server](#)

Use the **Users** screen on the **SSL VPN-Plus** tab to add accounts for your remote users to the local authentication server for the NSX Data Center for vSphere edge gateway SSL VPN service.

## 7 Add an SSL VPN-Plus Client Installation Package

Use the Installation Packages screen on the **SSL VPN-Plus** tab to create named installation packages of the SSL VPN-Plus client for the remote users.

## 8 Edit SSL VPN-Plus Client Configuration

Use the **Client Configuration** screen on the **SSL VPN-Plus** tab to customize the way the SSL VPN client tunnel responds when the remote user logs in to SSL VPN.

## 9 Customize the General SSL VPN-Plus Settings for an NSX Data Center for vSphere Edge Gateway

By default, the system sets some SSL VPN-Plus settings on an edge gateway in your VMware Cloud Director environment. You can use the **General Settings** screen on the **SSL VPN-Plus** tab in the VMware Cloud Director tenant portal to customize these settings.

### Navigate to the SSL-VPN Plus Screen

You can navigate to the SSL-VPN Plus screen to begin configuring the SSL-VPN Plus service for an NSX Data Center for vSphere edge gateway.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **SSL VPN-Plus** tab.

#### What to do next

On the **General** screen, configure the default SSL VPN-Plus settings. See [Customize the General SSL VPN-Plus Settings for an NSX Data Center for vSphere Edge Gateway](#).

### Configure SSL VPN Server Settings

These server settings configure the SSL VPN server, such as the IP address and port the service listens on, the cipher list of the service, and its service certificate. When connecting to the NSX Data Center for vSphere edge gateway, remote users specify the same IP address and port you set in these server settings.

If your edge gateway is configured with multiple, overlay IP address networks on its external interface, the IP address you select for the SSL VPN server can be different than the default external interface of the edge gateway.

While configuring the SSL VPN server settings, you must choose which encryption algorithms to use for the SSL VPN tunnel. You can choose one or more ciphers. Carefully choose the ciphers according to the strengths and weaknesses of your selections.

By default, the system uses the default, self-signed certificate that the system generates for each edge gateway as the default server identity certificate for the SSL VPN tunnel. Instead of this default, you can choose to use a digital certificate that you have added to the system on the **Certificates** screen.

#### Prerequisites

- Verify that you have met the prerequisites described in [Configure SSL VPN-Plus](#).
- If you choose to use a service certificate different than the default one, import the required certificate into the system. See [Add a Service Certificate to the Edge Gateway](#).
- [Navigate to the SSL-VPN Plus Screen](#).

#### Procedure

- 1 On the **SSL VPN-Plus** screen, click **Server Settings**.
- 2 Click **Enabled**.
- 3 Select an IP address from the drop-down menu.
- 4 (Optional) Enter a TCP port number.

The TCP port number is used by the SSL client installation package. By default, the system uses port 443, which is the default port for HTTPS/SSL traffic. Even though a port number is required, you can set any TCP port for communications.

---

**Note** The SSL VPN client requires the IP address and port configured here to be reachable from the client systems of your remote users. If you change the port number from the default, ensure that the IP address and port combination are reachable from the systems of your intended users.

---

- 5 Select an encryption method from the cipher list.
- 6 Configure the service Syslog logging policy.  
Logging is activated by default. You can change the level of messages to log or deactivate logging.
- 7 (Optional) If you want to use a service certificate instead of the default system-generated self-signed certificate, click **Change server certificate**, selection a certificate, and click **OK**.
- 8 Click **Save changes**.

#### What to do next

---

**Note** The edge gateway IP address and the TCP port number you set must be reachable by your remote users. Add an edge gateway firewall rule that allows access to the SSL VPN-Plus IP address and port configured in this procedure. See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#).

---

Add an IP pool so that remote users are assigned IP addresses when they connect using SSL VPN-Plus. See [Create an IP Pool for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#).

### Create an IP Pool for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway

The remote users are assigned virtual IP addresses from the static IP pools that you configure using the **IP Pools** screen on the **SSL VPN-Plus** tab.

Each IP pool added in this screen results in an IP address subnet configured on the edge gateway. The IP address ranges used in these IP pools must be different from all other networks configured on the edge gateway.

**Note** SSL VPN assigns IP addresses to the remote users from the IP pools based on the order the IP pools appear in the on-screen table. After you add the IP pools to the on-screen table, you can adjust their positions in the table using the up and down arrows.

#### Prerequisites

- [Navigate to the SSL-VPN Plus Screen.](#)
- [Configure SSL VPN Server Settings.](#)

#### Procedure

- 1 On the **SSL VPN-Plus** tab, click **IP Pools**.
- 2 Click the **Create** () button.
- 3 Configure the IP pool settings.

Option	Action
<b>IP Range</b>	Enter an IP address range for this IP pool, such as <b>127.0.0.1-127.0.0.9</b> . These IP addresses will be assigned to VPN clients when they authenticate and connect to the SSL VPN tunnel.
<b>Netmask</b>	Enter the netmask of the IP pool, such as <b>255.255.255.0</b> .
<b>Gateway</b>	Enter the IP address that you want the edge gateway to create and assign as the gateway address for this IP pool. When the IP pool is created, a virtual adapter is created on the edge gateway virtual machine and this IP address is configured on that virtual interface. This IP address can be any IP within the subnet that is not also in the range in the <b>IP Range</b> field.
<b>Description</b>	(Optional) Enter a description for this IP pool.
<b>Status</b>	Select whether to activate or deactivate this IP pool.
<b>Primary DNS</b>	(Optional) Enter the name of the primary DNS server that will be used for name resolution for these virtual IP addresses.
<b>Secondary DNS</b>	(Optional) Enter the name of the secondary DNS server to use.

Option	Action
DNS Suffix	(Optional) Enter the DNS suffix for the domain the client systems are hosted on, for domain-based host name resolution.
WINS Server	(Optional) Enter the WINS server address for the needs of your organization.

#### 4 Click **Keep**.

#### Results

The IP pool configuration is added to the on-screen table.

#### What to do next

Add private networks that you want accessible to your remote users connecting with SSL VPN-Plus. See [Add a Private Network for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#).

#### Add a Private Network for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway

Use the Private Networks screen on the **SSL VPN-Plus** tab to configure the private networks. The private networks are the ones you want the VPN clients to have access to, when the remote users connect using their VPN clients and the SSL VPN tunnel. The activated private networks will be installed in the routing table of the VPN client.

The private networks is a list of all reachable IP networks behind the edge gateway that you want to encrypt traffic for a VPN client, or exclude from encrypting. Each private network that requires access through an SSL VPN tunnel must be added as a separate entry. You can use route summarization techniques to limit the number of entries.

- SSL VPN-Plus allows remote users to access private networks based on the top-down order the IP pools appear in the on-screen table. After you add the private networks to the on-screen table, you can adjust their positions in the table using the up and down arrows.
- If you select to activate TCP optimization for a private network, some applications such as FTP in active mode might not work within that subnet. To add an FTP server configured in active mode, you must add another private network for that FTP server and deactivate TCP optimization for that private network. Also, the private network for that FTP server must be activated and appear in the on-screen table above the TCP-optimized private network.

#### Prerequisites

- [Navigate to the SSL-VPN Plus Screen](#).
- [Create an IP Pool for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#).

#### Procedure

- 1 On the **SSL VPN-Plus** tab, click **Private Networks**.

- 2 Click the **Add** () button.
- 3 Configure the private network settings.

Option	Action
<b>Network</b>	Type the private network IP address in a CIDR format, such as <b>192169.1.0/24</b> .
<b>Description</b>	(Optional) Type a description for the network.
<b>Send Traffic</b>	<p>Specify how you want the VPN client to send the private network and Internet traffic.</p> <ul style="list-style-type: none"> <li>■ <b>Over Tunnel</b> <p>The VPN client sends the private network and Internet traffic over the SSL VPN-Plus activated edge gateway.</p> </li> <li>■ <b>Bypass Tunnel</b> <p>The VPN client bypasses the edge gateway and sends the traffic directly to the private server.</p> </li> </ul>
<b>Enable TCP Optimization</b>	<p>(Optional) To best optimize the Internet speed, when you select <b>Over Tunnel</b> for sending the traffic, you must also select <b>Enable TCP Optimization</b>. Selecting this option enhances the performance of TCP packets within the VPN tunnel but does not improve performance of UDP traffic.</p> <p>Conventional full-access SSL VPNs tunnel sends TCP/IP data in a second TCP/IP stack for encryption over the Internet. This conventional method encapsulates application layer data in two separate TCP streams. When packet loss occurs, which can happen even under optimal Internet conditions, a performance degradation effect called TCP-over-TCP meltdown occurs. In TCP-over-TCP meltdown, two TCP instruments correct the same single packet of IP data, undermining network throughput and causing connection timeouts. Selecting <b>Enable TCP Optimization</b> eliminates the risk of this TCP-over-TCP problem occurring.</p> <p><b>Note</b> When you activate TCP optimization:</p> <ul style="list-style-type: none"> <li>■ You must enter the port numbers for which to optimize the Internet traffic.</li> <li>■ The SSL VPN server opens the TCP connection on behalf of the VPN client. When the SSL VPN server opens the TCP connection, the first automatically generated edge firewall rule is applied, which allows all connections opened from the edge gateway to get passed. Traffic that is not optimized is evaluated by the regular edge firewall rules. The default generated TCP rule is to allow any connections.</li> </ul>
<b>Ports</b>	<p>When you select <b>Over Tunnel</b>, type a range of port numbers that you want opened for the remote user to access the internal servers, such as <b>20–21</b> for FTP traffic and <b>80–81</b> for HTTP traffic.</p> <p>To give unrestricted access to users, leave the field blank.</p>
<b>Status</b>	Activate or deactivate the private network.

- 4 Click **Keep**.
- 5 Click **Save changes** to save the configuration to the system.

## What to do next

Add an authentication server. See [Configure an Authentication Service for SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#).

---

**Important** Add the corresponding firewall rules to allow network traffic to the private networks you have added in this screen. See [Add an NSX Data Center for vSphere Edge Gateway Firewall Rule](#).

---

## Configure an Authentication Service for SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway

Use the **Authentication** screen on the **SSL VPN-Plus** tab to set up a local authentication server for the edge gateway SSL VPN service and optionally enable client certificate authentication. This authentication server is used to authenticate the connecting users. All users configured in the local authentication server will be authenticated.

You can have only one local SSL VPN-Plus authentication server configured on the edge gateway. If you click **+ LOCAL** and specify additional authentication servers, an error message is displayed when you try to save the configuration.

The maximum time to authenticate over SSL VPN is three (3) minutes. This maximum is determined by the non-authentication timeout, which is 3 minutes by default and is not configurable. As a result, if you have multiple authentication servers in chain authorization and user authentication takes more than 3 minutes, the user will not be authenticated.

### Prerequisites

- [Navigate to the SSL-VPN Plus Screen](#).
- [Add a Private Network for Use with SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#).
- If you intend to enable client certificate authentication, verify that a CA certificate has been added to the edge gateway. See [Add a CA Certificate to the Edge Gateway for SSL Certificate Trust Verification](#).

### Procedure

- 1 Click the **SSL VPN-Plus** tab and **Authentication**.
- 2 Click **Local**.

### 3 Configure the authentication server settings.

- a (Optional) Enable and configure the password policy.

Option	Description
<b>Enable password policy</b>	Turn on enforcement of the password policy settings you configure here.
<b>Password Length</b>	Enter the minimum and maximum allowed number of characters for password length.
<b>Minimum no. of alphabets</b>	(Optional) Type the minimum number of alphabetic characters, that are required in the password.
<b>Minimum no. of digits</b>	(Optional) Type the minimum number of numeric characters, that are required in the password.
<b>Minimum no. of special characters</b>	(Optional) Type the minimum number of special characters, such as ampersand (&), hash tag (#), percent sign (%) and so on, that are required in the password.
<b>Password should not contain user ID</b>	(Optional) Enable to enforce that the password must not contain the user ID.
<b>Password expires in</b>	(Optional) Type the maximum number of days that a password can exist before the user must change it.
<b>Expiry notification in</b>	(Optional) Type the number of days prior to the <b>Password expires in</b> value at which the user is notified the password is about to expire.

- b (Optional) Enable and configure the account lockout policy.

Option	Description
<b>Enable account lockout policy</b>	Turn on enforcement of the account lockout policy settings you configure here.
<b>Retry Count</b>	Enter the number of times a user can try to access their account.
<b>Retry Duration</b>	Enter the time period in minutes in which the user account gets locked on unsuccessful login attempts. For example, if you specify the <b>Retry Count</b> as 5 and <b>Retry Duration</b> as 1 minute, the account of the user is locked after 5 unsuccessful login attempts within 1 minute.
<b>Lockout Duration</b>	Enter the time period for which the user account remains locked. After this time has elapsed, the account is automatically unlocked.

- c In the Status section, enable this authentication server.

- d (Optional) Configure secondary authentication.

Options	Description
<b>Use this server for secondary authentication</b>	(Optional) Specify whether to use the server as the second level of authentication.
<b>Terminate session if authentication fails</b>	(Optional) Specify whether to end the VPN session when authentication fails.

- e Click **Keep**.

- 4 (Optional) To enable client certification authentication, click **Change certificate**, then turn on the enablement toggle, select the CA certificate to use, and click **OK**.

### What to do next

Add local users to the local authentication server so that they can connect with SSL VPN-Plus. See [Add SSL VPN-Plus Users to the Local SSL VPN-Plus Authentication Server](#).

Create an installation package containing the SSL Client so remote users can install it on their local systems. See [Add an SSL VPN-Plus Client Installation Package](#).

### Add SSL VPN-Plus Users to the Local SSL VPN-Plus Authentication Server

Use the **Users** screen on the **SSL VPN-Plus** tab to add accounts for your remote users to the local authentication server for the NSX Data Center for vSphere edge gateway SSL VPN service.

**Note** If a local authentication server is not already configured, adding a user on the **Users** screen automatically adds a local authentication server with default values. You can then use the edit button on the **Authentication** screen to view and edit the default values. For information about using the **Authentication** screen, see [Configure an Authentication Service for SSL VPN-Plus on an NSX Data Center for vSphere Edge Gateway](#).

### Prerequisites

[Navigate to the SSL-VPN Plus Screen](#).

### Procedure

- 1 On the **SSL VPN-Plus** tab, click **Users**.
- 2 Click the **Create** () button.
- 3 Configure the following options for the user.

Option	Description
<b>User ID</b>	Enter the user ID.
<b>Password</b>	Enter a password for the user.
<b>Retype Password</b>	Reenter the password.
<b>First name</b>	(Optional) Enter the first name of the user.

Option	Description
Last name	(Optional) Enter the last name of the user.
Description	(Optional) Enter a description for the user.
Enabled	Specify whether the user is activated or deactivated.
Password never expires	(Optional) Specify whether to keep the same password for this user forever.
Allow change password	(Optional) Specify whether to let the user change the password.
Change password on next login	(Optional) Specify whether you want this user to change the password the next time the user logs in.

- 4 Click **Keep**.
- 5 Repeat the steps to add additional users.

#### What to do next

Add local users to the local authentication server so that they can connect with SSL VPN-Plus. See [Add SSL VPN-Plus Users to the Local SSL VPN-Plus Authentication Server](#).

Create an installation package containing the SSL Client so the remote users can install it on their local systems. See [Add an SSL VPN-Plus Client Installation Package](#).

#### Add an SSL VPN-Plus Client Installation Package

Use the Installation Packages screen on the **SSL VPN-Plus** tab to create named installation packages of the SSL VPN-Plus client for the remote users.

You can add an SSL VPN-Plus client installation package to the NSX Data Center for vSphere edge gateway. New users are prompted to download and install this package when they log in to use the VPN connection for the first time. When added, these client installation packages are then downloadable from the FQDN of the edge gateway's public interface.

You can create installation packages that run on Windows, Linux, and Mac operating systems. If you require different installation parameters per SSL VPN client, create an installation package for each configuration.

#### Prerequisites

[Navigate to the SSL-VPN Plus Screen](#)

#### Procedure

- 1 On the **SSL VPN-Plus** tab in the tenant portal, click **Installation Packages**.
- 2 Click the **Add** () button.

### 3 Configure the installation package settings.

Option	Description
<b>Profile Name</b>	Enter a profile name for this installation package. This name is displayed to the remote user to identify this SSL VPN connection to the edge gateway.
<b>Gateway</b>	Enter the IP address or FQDN of the edge gateway public interface. The IP address or FQDN that you enter is bound to the SSL VPN client. When the client is installed on the local system of the remote user, this IP address or FQDN is displayed on that SSL VPN client. To bind additional edge gateway uplink interfaces to this SSL VPN client, click the <b>Add</b> (  ) button to add rows and type in their interface IP addresses or FQDNs, and ports.
<b>Port</b>	(Optional) To modify the port value from the displayed default, double-click the value and enter a new value.
<b>Windows</b> <b>Linux</b> <b>Mac</b>	Select the operating systems for which you want to create the installation packages.
<b>Description</b>	(Optional) Type a description for the user.
<b>Enabled</b>	Specify whether this package is activated or deactivated.

### 4 Select the installation parameters for Windows.

Option	Description
<b>Start client on logon</b>	Starts the SSL VPN client when the remote user logs in to their local system.
<b>Allow remember password</b>	Enables the client to remember the user password.
<b>Enable silent mode installation</b>	Hides installation commands from remote users.
<b>Hide SSL client network adapter</b>	Hides the VMware SSL VPN-Plus Adapter which is installed on the computer of the remote user, together with the SSL VPN client installation package.
<b>Hide client system tray icon</b>	Hides the SSL VPN tray icon which indicates whether the VPN connection is active or not.
<b>Create desktop icon</b>	Creates an icon on the user desktop to invoke the SSL client.
<b>Enable silent mode operation</b>	Hides the window that indicates that installation is complete.
<b>Server security certificate validation</b>	The SSL VPN client validates the SSL VPN server certificate before establishing the secure connection.

### 5 Click **Keep**.

#### What to do next

Edit the client configuration. See [Edit SSL VPN-Plus Client Configuration](#).

## Edit SSL VPN-Plus Client Configuration

Use the **Client Configuration** screen on the **SSL VPN-Plus** tab to customize the way the SSL VPN client tunnel responds when the remote user logs in to SSL VPN.

### Prerequisites

[Navigate to the SSL-VPN Plus Screen](#)

### Procedure

- 1 On the **SSL VPN-Plus** tab, click **Client Configuration**.
- 2 Select the **Tunneling mode**.
  - In split tunnel mode, only the VPN traffic flows through the edge gateway.
  - In full tunnel mode, the edge gateway becomes the default gateway for the remote user and all traffic, such as VPN, local, and Internet, flows through the edge gateway.
- 3 If you select full tunnel mode, enter the IP address for the default gateway used by the clients of the remote users and, optionally, select whether to exclude local subnet traffic from flowing through the VPN tunnel.
- 4 (Optional) Deactivate auto reconnect.
 

**Enable auto reconnect** is activated by default. If auto reconnect is activated, the SSL VPN client automatically reconnects users when they get disconnected.
- 5 (Optional) Optionally enable the ability for the client to notify remote users when a client upgrade is available.
 

This option is deactivated by default. If you activate this option, remote users can choose to install the upgrade.
- 6 Click **Save changes**.

## Customize the General SSL VPN-Plus Settings for an NSX Data Center for vSphere Edge Gateway

By default, the system sets some SSL VPN-Plus settings on an edge gateway in your VMware Cloud Director environment. You can use the **General Settings** screen on the **SSL VPN-Plus** tab in the VMware Cloud Director tenant portal to customize these settings.

### Prerequisites

[Navigate to the SSL-VPN Plus Screen](#).

### Procedure

- 1 On the **SSL VPN-Plus** tab, click **General Settings**.

## 2 Edit the general settings as required for the needs of your organization.

Option	Description
<b>Prevent multiple logon using same username</b>	Turn on to restrict a remote user to having only one active login session under the same user name.
<b>Compression</b>	Turn on to enable TCP-based intelligent data compression and improve data transfer speed.
<b>Enable Logging</b>	Turn on to maintain a log of the traffic that passes through the SSL VPN gateway. Logging is enabled by default.
<b>Force virtual keyboard</b>	Turn on to require remote users to use a virtual (on-screen) keyboard only to enter login information.
<b>Randomize keys of virtual keyboard</b>	Turn on to have the virtual keyboard use a randomized key layout.
<b>Session idle timeout</b>	Enter the session idle timeout in minutes. If there is no activity in a user session for the specified time period, the system disconnects the user session. The system default is 10 minutes.
<b>User notification</b>	Type the message to be displayed to remote users after they log in.
<b>Enable public URL access</b>	Turn on to allow remote users to access sites that are not explicitly configured by you for remote user access.
<b>Enable forced timeout</b>	Turn on to have the system disconnect remote users after the time period that you specify in the <b>Forced timeout</b> field is over.
<b>Forced timeout</b>	Type the timeout period in minutes. This field is displayed when <b>Enable forced timeout</b> toggle is turned on.

## 3 Click **Save changes**.

### Configure IPsec VPN

The NSX Data Center for vSphere edge gateways in a VMware Cloud Director environment support site-to-site Internet Protocol Security (IPsec) to secure VPN tunnels between organization virtual data center networks or between an organization virtual data center network and an external IP address. You can configure the IPsec VPN service on an edge gateway.

Setting up an IPsec VPN connection from a remote network to your organization virtual data center is the most common scenario. The NSX software provides an edge gateway IPsec VPN capabilities, including support for certificate authentication, preshared key mode, and IP unicast traffic between itself and remote VPN routers. You can also configure multiple subnets to connect through IPsec tunnels to the internal network behind an edge gateway. When you configure multiple subnets to connect through IPsec tunnels to the internal network, those subnets and the internal network behind the edge gateway must not have address ranges that overlap.

**Note** If the local and remote peer across an IPsec tunnel have overlapping IP addresses, traffic forwarding across the tunnel might not be consistent depending on whether local connected routes and auto-plumbed routes exist.

The following IPsec VPN algorithms are supported:

- AES (AES128-CBC)
- AES256 (AES256-CBC)
- Triple DES (3DES192-CBC)
- AES-GCM (AES128-GCM)
- DH-2 (Diffie-Hellman group 2)
- DH-5 (Diffie-Hellman group 5)
- DH-14 (Diffie-Hellman group 14)

---

**Note** Dynamic routing protocols are not supported with IPsec VPN. When you configure an IPsec VPN tunnel between an edge gateway of the organization virtual data center and a physical gateway VPN at a remote site, you cannot configure dynamic routing for that connection. The IP address of that remote site cannot be learned by dynamic routing on the edge gateway uplink.

---

As described in the *IPsec VPN Overview* topic in the *NSX Administration Guide*, the maximum number of tunnels supported on an edge gateway is determined by its configured size: compact, large, x-large, quad large.

To view the size of your edge gateway configuration, navigate to the edge gateway and click the edge gateway name.

Configuring IPsec VPN on an edge gateway is a multi-step process.

---

**Note** If a firewall is between the tunnel endpoints, after you configure the IPsec VPN service, update the firewall rules to allow the following IP protocols and UDP ports:

- IP Protocol ID 50 (ESP)
  - IP Protocol ID 51 (AH)
  - UDP Port 500 (IKE)
  - UDP Port 4500
- 

## Procedure

### 1 [Navigate to the IPsec VPN Screen](#)

In the **IPsec VPN** screen, you can begin configuring the IPsec VPN service for an NSX Data Center for vSphere edge gateway.

### 2 [Configure the IPsec VPN Site Connections for the NSX Data Center for vSphere Edge Gateway](#)

Use the **IPsec VPN Sites** screen in the VMware Cloud Director tenant portal to configure settings needed to create an IPsec VPN connection between your organization virtual data center and another site using the edge gateway IPsec VPN capabilities.

### 3 Enable the IPsec VPN Service on an NSX Data Center for vSphere Edge Gateway

When at least one IPsec VPN connection is configured, you can enable the IPsec VPN service on the edge gateway.

### 4 Specify Global IPsec VPN Settings

Use the **Global Configuration** screen to configure IPsec VPN authentication settings at an edge gateway level. On this screen, you can set a global pre-shared key and enable certification authentication.

## Navigate to the IPsec VPN Screen

In the **IPsec VPN** screen, you can begin configuring the IPsec VPN service for an NSX Data Center for vSphere edge gateway.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **VPN > IPsec VPN**.

### What to do next

Use the **IPsec VPN Sites** screen to configure an IPsec VPN connection. At least one connection must be configured before you can enable the IPsec VPN service on the edge gateway. See [Configure the IPsec VPN Site Connections for the NSX Data Center for vSphere Edge Gateway](#).

## Configure the IPsec VPN Site Connections for the NSX Data Center for vSphere Edge Gateway

Use the **IPsec VPN Sites** screen in the VMware Cloud Director tenant portal to configure settings needed to create an IPsec VPN connection between your organization virtual data center and another site using the edge gateway IPsec VPN capabilities.

When you configure an IPsec VPN connection between sites, you configure the connection from the point of view of your current location. Setting up the connection requires that you understand the concepts in the context of the VMware Cloud Director environment so that you configure the VPN connection correctly.

- The local and peer subnets specify the networks to which the VPN connects. When you specify these subnets in the configurations for IPsec VPN sites, enter a network range and not a specific IP address. Use CIDR format, such as **192.168.99.0/24**.
- The peer ID is an identifier that uniquely identifies the remote device that terminates the VPN connection, typically its public IP address. For peers using certificate authentication, this ID must be the distinguished name set in the peer certificate. For PSK peers, this ID can be any string. An NSX best practice is to use the public IP address of the remote device or FQDN as the peer ID. If the peer IP address is from another organization virtual data center network, you enter the native IP address of the peer. If NAT is configured for the peer, you enter the peer's private IP address.

- The peer endpoint specifies the public IP address of the remote device to which you are connecting. The peer endpoint might be a different address from the peer ID if the peer's gateway is not directly accessible from the Internet, but connects through another device. If NAT is configured for the peer, you enter the public IP address that the devices uses for NAT.
- The local ID specifies the public IP address of the edge gateway of the organization virtual data center. You can enter an IP address or hostname along with the edge gateway firewall.
- The local endpoint specifies the network in your organization virtual data center on which the edge gateway transmits. Typically the external network of the edge gateway is the local endpoint.

### Prerequisites

- [Navigate to the IPsec VPN Screen.](#)
- [Configure IPsec VPN.](#)
- If you intend to use a global certificate as the authentication method, verify that certificate authentication is enabled on the **Global Configuration** screen. See [Specify Global IPsec VPN Settings.](#)

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 On the **IPsec VPN** tab, click **IPsec VPN Sites**.
- 3 Click the **Add** () button.
- 4 Configure the IPsec VPN connection settings.

Option	Action
<b>Enabled</b>	Enable this connection between the two VPN endpoints.
<b>Enable perfect forward secrecy (PFS)</b>	<p>Enable this option to have the system generate unique public keys for all IPsec VPN sessions your users initiate.</p> <p>Enabling PFS ensures that the system does not create a link between the edge gateway private key and each session key.</p> <p>The compromise of a session key will not affect data other than the data exchanged in the specific session protected by that particular key. Compromise of the server's private key cannot be used to decrypt archived sessions or future sessions.</p> <p>When PFS is enabled, IPsec VPN connections to this edge gateway experience a slight processing overhead.</p> <p><b>Important</b> The unique session keys must not be used to derive any additional keys. Also, both sides of the IPsec VPN tunnel must support PFS for it to work.</p>

Option	Action
<b>Name</b>	(Optional) Enter a name for the connection.
<b>Local ID</b>	<p>Enter the external IP address of the edge gateway instance, which is the public IP address of the edge gateway.</p> <p>The IP address is the one used for the peer ID in the IPsec VPN configuration on the remote site.</p>
<b>Local Endpoint</b>	<p>Enter the network that is the local endpoint for this connection.</p> <p>The local endpoint specifies the network in your organization virtual data center on which the edge gateway transmits. Typically, the external network is the local endpoint.</p> <p>If you add an IP-to-IP tunnel using a pre-shared key, the local ID and local endpoint IP can be the same.</p>
<b>Local Subnets</b>	<p>Enter the networks to share between the sites and use a comma as a separator to enter multiple subnets.</p> <p>Enter a network range (not a specific IP address) by entering the IP address using CIDR format. For example, <b>192 . 168 . 99 . 0/24</b>.</p>
<b>Peer ID</b>	<p>Enter a peer ID to uniquely identify the peer site.</p> <p>The peer ID is an identifier that uniquely identifies the remote device that terminates the VPN connection, typically its public IP address.</p> <p>For peers using certificate authentication, the ID must be the distinguished name in the peer's certificate. For PSK peers, this ID can be any string. An NSX best practice is to use the remote device's public IP address or FQDN as the peer ID.</p> <p>If the peer IP address is from another organization virtual data center network, you enter the native IP address of the peer. If NAT is configured for the peer, you enter the peer's private IP address.</p>
<b>Peer Endpoint</b>	<p>Enter the IP address or FQDN of the peer site, which is the public-facing address of the remote device to which you are connecting.</p> <p><b>Note</b> When NAT is configured for the peer, enter the public IP address that the device uses for NAT.</p>
<b>Peer Subnets</b>	<p>Enter the remote network to which the VPN connects and use a comma as a separator to enter multiple subnets.</p> <p>Enter a network range (not a specific IP address) by entering the IP address using CIDR format. For example, <b>192 . 168 . 99 . 0/24</b>.</p>
<b>Encryption Algorithm</b>	<p>Select the encryption algorithm type from the drop-down menu.</p> <p><b>Note</b> The encryption type you select must match the encryption type configured on the remote site VPN device.</p>

Option	Action
Authentication	<p>Select an authentication. The options are:</p> <ul style="list-style-type: none"> <li>■ <b>PSK</b></li> </ul> <p>Pre Shared Key (PSK) specifies that the secret key shared between the edge gateway and the peer site is to be used for authentication.</p> <ul style="list-style-type: none"> <li>■ <b>Certificate</b></li> </ul> <p>Certificate authentication specifies that the certificate defined at the global level is to be used for authentication. This option is not available unless you have configured the global certificate on the <b>IPsec VPN</b> tab's <b>Global Configuration</b> screen.</p>
Change Shared Key	(Optional) When you are updating the settings of an existing connection, you can turn on this option on to make the <b>Pre-Shared Key</b> field available so that you can update the shared key.
Pre-Shared Key	<p>If you selected <b>PSK</b> as the authentication type, type an alphanumeric secret string which can be a string with a maximum length of 128 bytes.</p> <p><b>Note</b> The shared key must match the key that is configured on the remote site VPN device. A best practice is to configure a shared key when anonymous sites will connect to the VPN service.</p>
Display Shared Key	(Optional) Enable this option to make the shared key visible in the screen.
Diffie-Hellman Group	<p>Select the cryptography scheme that allows the peer site and this edge gateway to establish a shared secret over an insecure communications channel.</p> <p><b>Note</b> The Diffie-Hellman Group must match what is configured on the remote site VPN device.</p>
Extension	<p>(Optional) Type one of the following options:</p> <ul style="list-style-type: none"> <li>■ <code>securelocaltrafficbyip=IPAddress</code> to redirect the edge gateway local traffic over the IPsec VPN tunnel.</li> </ul> <p>This is the default value.</p> <ul style="list-style-type: none"> <li>■ <code>passthroughSubnets=PeerSubnet/IPAddress</code> to support overlapping subnets.</li> </ul>

5 Click **Keep**.

6 Click **Save changes**.

#### What to do next

Configure the connection for the remote site. You must configure the IPsec VPN connection on both sides of the connection: your organization virtual data center and the peer site.

Enable the IPsec VPN service on this edge gateway. When at least one IPsec VPN connection is configured, you can enable the service. See [Enable the IPsec VPN Service on an NSX Data Center for vSphere Edge Gateway](#).

## Enable the IPsec VPN Service on an NSX Data Center for vSphere Edge Gateway

When at least one IPsec VPN connection is configured, you can enable the IPsec VPN service on the edge gateway.

### Prerequisites

- [Navigate to the IPsec VPN Screen.](#)
- Verify that at least one IPsec VPN connection is configured for this edge gateway. See the steps described in [Configure the IPsec VPN Site Connections for the NSX Data Center for vSphere Edge Gateway.](#)

### Procedure

- 1 On the **IPsec VPN** tab, click **Activation Status**.
- 2 Click **IPsec VPN Service Status** to enable the IPsec VPN service.
- 3 Click **Save changes**.

### Results

The edge gateway IPsec VPN service is active.

## Specify Global IPsec VPN Settings

Use the **Global Configuration** screen to configure IPsec VPN authentication settings at an edge gateway level. On this screen, you can set a global pre-shared key and enable certification authentication.

A global pre-shared key is used for those sites whose peer endpoint is set to **any**.

### Prerequisites

- If you intend to enable certificate authentication, verify that you have at least one service certificate and corresponding CA-signed certificates in the **Certificates** screen. Self-signed certificates cannot be used for IPsec VPNs. See [Add a Service Certificate to the Edge Gateway.](#)
- [Navigate to the IPsec VPN Screen.](#)

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 On the **IPsec VPN** tab, click **Global Configuration**.

### 3 (Optional) Set a global pre-shared key:

- a Enable the **Change Shared Key** option.
- b Enter a pre-shared key.

The global pre-shared key (PSK) is shared by all the sites whose peer endpoint is set to any. If a global PSK is already set, changing the PSK to an empty value and saving it has no effect on the existing setting.

- c (Optional) Optionally enable **Display Shared Key** to make the pre-shared key visible.
- d Click **Save changes**.

### 4 Configure certification authentication:

- a Turn on **Enable Certificate Authentication**.
- b Select the appropriate service certificates, CA certificates, and CRLs.
- c Click **Save changes**.

#### What to do next

You can optionally enable logging for the IPsec VPN service of the edge gateway. See [Statistics and Logs for an NSX Data Center for vSphere Edge Gateway](#).

## Configure L2 VPN

The NSX Data Center for vSphere edge gateways in a VMware Cloud Director environment support L2 VPN. With L2 VPN, you can extend your organization virtual data center by enabling virtual machines to maintain network connectivity while retaining the same IP address across geographical boundaries. You can configure the L2 VPN service on an edge gateway.

NSX Data Center for vSphere provides the L2 VPN capabilities of an edge gateway. With L2 VPN, you can configure a tunnel between two sites. Virtual machines remain on the same subnet despite being moved between these sites, which enables you to extend your organization virtual data center by stretching its network using L2 VPN. An edge gateway at one site can provide all services to virtual machines on the other site.

To create the L2 VPN tunnel, you configure an L2 VPN server and L2 VPN client. As described in the *NSX Administration Guide*, the L2 VPN server is the destination edge gateway and the L2 VPN client is the source edge gateway. After configuring the L2 VPN settings on each edge gateway, you must then enable the L2 VPN service on both the server and the client.

---

**Note** A routed organization virtual data center network created as a subinterface must exist on the edge gateways.

---

#### Navigate to the L2 VPN Screen

To begin configuring the L2 VPN service for an NSX Data Center for vSphere edge gateway, you must navigate to the **L2 VPN** screen.

**Procedure**

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Navigate to **VPN > L2 VPN**.

**What to do next**

Configure the L2 VPN server. See [Configure the NSX Data Center for vSphere Edge Gateway as an L2 VPN Server](#).

**Configure the NSX Data Center for vSphere Edge Gateway as an L2 VPN Server**

The L2 VPN server is the destination NSX edge to which the L2 VPN client is going to connect.

As described in the *NSX Administration Guide*, you can connect multiple peer sites to this L2 VPN server.

---

**Note** Changing site configuration settings causes the edge gateway to disconnect and reconnect all existing connections.

---

**Prerequisites**

- Verify that the edge gateway has a routed organization virtual data center network that is configured as a subinterface on the edge gateway.
- [Navigate to the L2 VPN Screen](#).
- If you want to bind a service certificate to the L2 VPN connection, verify that the server certificate has already been uploaded to the edge gateway. See [Add a Service Certificate to the Edge Gateway](#).
- You must have the listener IP of the server, listener port, encryption algorithm, and at least one peer site configured before you can enable the L2 VPN service.

**Procedure**

- 1 On the **L2 VPN** tab, select **Server** for the L2 VPN mode.
- 2 On the **Server Global** tab, configure the L2 VPN server's global configuration details.

Option	Action
Listener IP	Select the primary or secondary IP address of an external interface of the edge gateway.
Listener Port	Edit the displayed value as appropriate for the needs of your organization. The default port for the L2 VPN service is 443.

Option	Action
<b>Encryption Algorithm</b>	Select the encryption algorithm for the communication between the server and the client.
<b>Service Certificate Details</b>	Click <b>Change server certificate</b> to select the certificate to be bound to the L2 VPN server. In the <b>Change Server Certificate</b> window, turn on <b>Validate Server Certificate</b> , select a server certificate from the list, and click <b>OK</b> .

- 3 To configure the peer sites, click the **Server Sites** tab.
- 4 Click the **Add** button.
- 5 Configure the settings for an L2 VPN peer site.

Option	Action
<b>Enabled</b>	Enable this peer site.
<b>Name</b>	Enter a unique name for the peer site.
<b>Description</b>	(Optional) Type a description.
<b>User ID</b> <b>Password</b> <b>Confirm Password</b>	Enter the user name and password with which the peer site is to be authenticated. User credentials on the peer site must be the same as the credentials on the client side.
<b>Stretched Interfaces</b>	Select at least one subinterface to be stretched with the client. The subinterfaces available for selection are those organization virtual data center networks configured as subinterfaces on the edge gateway.
<b>Egress Optimization Gateway Address</b>	(Optional) If the default gateway for virtual machines is the same across the two sites, enter the gateway IP addresses of the subinterfaces for which you want the traffic locally routed or blocked over the L2 VPN tunnel.

- 6 Click **Keep**.
- 7 Click **Save changes**.

#### What to do next

Enable the L2 VPN service on this edge gateway. See [Enable the L2 VPN Service on an NSX Data Center for vSphere Edge Gateway](#).

#### Configure the NSX Data Center for vSphere Edge Gateway as an L2 VPN Client

The L2 VPN client is the source NSX edge that initiates communication with the destination NSX edge, the L2 VPN server.

#### Prerequisites

- [Navigate to the L2 VPN Screen](#).

- If this L2 VPN client is connecting to an L2 VPN server that uses a server certificate, verify that the corresponding CA certificate is uploaded to the edge gateway to enable server certificate validation for this L2 VPN client. See [Add a CA Certificate to the Edge Gateway for SSL Certificate Trust Verification](#).

#### Procedure

- 1 On the **L2 VPN** tab, select **Client** for the L2 VPN mode.
- 2 On the **Client Global** tab, configure the global configuration details of the L2 VPN client.

Option	Description
<b>Server Address</b>	Enter the IP address of the L2 VPN server to which this client is to be connected.
<b>Server Port</b>	Enter the L2 VPN server port to which the client should connect. The default port is 443.
<b>Encryption Algorithm</b>	Select the encryption algorithm for communicating with the server.
<b>Stretched Interfaces</b>	Select the subinterfaces to be stretched to the server. The subinterfaces available to select are the organization virtual data center networks configured as subinterfaces on the edge gateway.
<b>Egress Optimization Gateway Address</b>	(Optional) If the default gateway for virtual machines is the same across the two sites, type the gateway IP addresses of the subinterfaces or the IP addresses to which traffic should not flow over the tunnel.
<b>User Details</b>	Enter the user ID and password for authentication with the server.

- 3 Click **Save changes**.
- 4 (Optional) To configure advanced options, click the **Client Advanced** tab.
- 5 If this L2 VPN client edge does not have direct access to the Internet, and must reach the L2 VPN server edge by using a proxy server, specify the proxy settings.

Option	Description
<b>Enable Secure Proxy</b>	Select to enable the secure proxy.
<b>Address</b>	Enter the proxy server IP address.
<b>Port</b>	Enter the proxy server port.
<b>User Name</b>	Enter the proxy server authentication credentials.
<b>Password</b>	

- 6 To enable server certification validation, click **Change CA certificate** and select the appropriate CA certificate.
- 7 Click **Save changes**.

#### What to do next

Enable the L2 VPN service on this edge gateway. See [Enable the L2 VPN Service on an NSX Data Center for vSphere Edge Gateway](#).

## Enable the L2 VPN Service on an NSX Data Center for vSphere Edge Gateway

When the required L2 VPN settings are configured, you can enable the L2 VPN service on the edge gateway.

---

**Note** If HA is already configured on this edge gateway, ensure that the edge gateway has more than one internal interface configured on it. If only a single interface exists and that has already been used by the HA capability, the L2 VPN configuration on the same internal interface fails.

---

### Prerequisites

- If this edge gateway is an L2 VPN server, the destination NSX edge, verify that the required L2 VPN server settings and at least one L2 VPN peer site are configured. See the steps described in [Configure the NSX Data Center for vSphere Edge Gateway as an L2 VPN Server](#).
- If this edge gateway is an L2 VPN client, the source NSX edge, verify that the L2 VPN client settings are configured. See the steps described in [Configure the NSX Data Center for vSphere Edge Gateway as an L2 VPN Client](#).
- [Navigate to the L2 VPN Screen](#).

### Procedure

- 1 On the **L2 VPN** tab, click the **Enable** toggle.
- 2 Click **Save changes**.

### Results

The L2 VPN service of the edge gateway becomes active.

### What to do next

Create NAT or firewall rules on the Internet-facing firewall side to enable the L2 VPN server to connect to the L2 VPN client.

## Remove the L2 VPN Service Configuration from an NSX Data Center for vSphere Edge Gateway

You can remove the existing L2 VPN service configuration of the edge gateway. This action also deactivates the L2 VPN service on the edge gateway.

### Prerequisites

[Navigate to the L2 VPN Screen](#)

### Procedure

- 1 Scroll down to the bottom of the L2 VPN screen, and click **Delete configuration**.
- 2 To confirm the deletion, click **OK**.

## Results

The L2 VPN service is deactivated and the configuration details are removed from the edge gateway.

## SSL Certificate Management on an NSX Data Center for vSphere Edge Gateway

The NSX Data Center for vSphere software in the VMware Cloud Director environment provides the ability to use Secure Sockets Layer (SSL) certificates with the SSL VPN-Plus and IPsec VPN tunnels you configure for your edge gateways.

The edge gateways in your VMware Cloud Director environment support self-signed certificates, certificates signed by a Certification Authority (CA), and certificates generated and signed by a CA. You can generate certificate signing requests (CSRs), import the certificates, manage the imported certificates, and create certificate revocation lists (CRLs).

### About Using Certificates with Your Organization Virtual Data Center

You can manage certificates for the following networking areas in your VMware Cloud Director organization virtual data center.

- IPsec VPN tunnels between an organization virtual data center network and a remote network.
- SSL VPN-Plus connections between remote users to private networks and web resources in your organization virtual data center.
- An L2 VPN tunnel between two NSX Data Center for vSphere edge gateways.
- The virtual servers and pools servers configured for load balancing in your organization virtual data center

### How to Use Client Certificates

You can create a client certificate through a CLI command or REST call. You can then distribute this certificate to your remote users, who can install the certificate on their web browser.

The main benefit of implementing client certificates is that a reference client certificate for each remote user can be stored and checked against the client certificate presented by the remote user. To prevent future connections from a certain user, you can delete the reference certificate from the security server list of client certificates. Deleting the certificate denies connections from that user.

### Generate a Certificate Signing Request for an Edge Gateway

Before you can order a signed certificate from a CA or create a self-signed certificate, you must generate a Certificate Signing Request (CSR) for your edge gateway.

A CSR is an encoded file that you need to generate on an NSX edge gateway which requires an SSL certificate. Using a CSR standardizes the way that companies send their public keys together with information that identifies their company names and domain names.

You generate a CSR with a matching private-key file that must remain on the edge gateway. The CSR contains the matching public key and other information such as the name, location, and domain name of your organization.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Certificates** tab.
- 3 On the **Certificates** tab, click **CSR**.
- 4 Configure the following options for the CSR:

Option	Description
<b>Common Name</b>	Enter the fully qualified domain name (FQDN) for the organization that you will be using the certificate for (for example, <code>www.example.com</code> ). Do not include the <code>http://</code> or <code>https://</code> prefixes in your common name.
<b>Organization Unit</b>	Use this field to differentiate between divisions within your VMware Cloud Director organization with which this certificate is associated. For example, Engineering or Sales.
<b>Organization Name</b>	Enter 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.
<b>Locality</b>	Enter the city or locality where your company is legally registered.
<b>State or Province Name</b>	Enter the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
<b>Country Code</b>	Enter the country name where your company is legally registered.
<b>Private Key Algorithm</b>	Type the key type, either RSA or DSA, for the certificate. RSA is typically used. The key type defines the encryption algorithm for communication between the hosts.  <b>Note</b> SSL VPN-Plus supports RSA certificates only.
<b>Key Size</b>	Enter the key size in bits. The minimum is 2048 bits.
<b>Description</b>	(Optional) Enter a description for the certificate.

- 5 Click **Keep**.

The system generates the CSR and adds a new entry with type CSR to the on-screen list.

#### Results

In the on-screen list, when you select an entry with type CSR, the CSR details are displayed in the screen. You can copy the displayed PEM formatted data of the CSR and submit it to a certificate authority (CA) to obtain a CA-signed certificate.

## What to do next

Use the CSR to create a service certificate using one of these two options:

- Transmit the CSR to a CA to obtain a CA-signed certificate. When the CA sends you the signed certificate, import the signed certificate into the system. See [Import the CA-Signed Certificate Corresponding to the CSR Generated for an Edge Gateway](#).
- Use the CSR to create a self-signed certificate. See [Configure a Self-Signed Service Certificate](#).

## Import the CA-Signed Certificate Corresponding to the CSR Generated for an Edge Gateway

After you generate a Certificate Signing Request (CSR) and obtain the CA-signed certificate based on that CSR, you can import the CA-signed certificate to use it by your edge gateway.

### Prerequisites

Verify that you obtained the CA-signed certificate that corresponds to the CSR. If the private key in the CA-signed certificate does not match the one for the selected CSR, the import process fails.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Certificates** tab.
- 3 Select the CSR in the on-screen table for which you are importing the CA-signed certificate.
- 4 Import the signed certificate.
  - a Click **Signed certificate generated for CSR**.
  - b Provide the PEM data of the CA-signed certificate.
    - If the data is in a PEM file on a system you can navigate to, click the **Upload** button to browse to the file and select it.
    - If you can copy and paste the PEM data, paste it into the **Signed Certificate (PEM format)** field.
 

Include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` lines.
  - c (Optional) Type a description.
  - d Click **Keep**.

---

**Note** If the private key in the CA-signed certificate does not match the one for the CSR you selected on the Certificates screen, the import process fails.

---

## Results

The CA-signed certificate with type Service Certificate appears in the on-screen list.

## What to do next

Attach the CA-signed certificate to your SSL VPN-Plus or IPsec VPN tunnels as required. See [Configure SSL VPN Server Settings](#) and [Specify Global IPsec VPN Settings](#).

## Configure a Self-Signed Service Certificate

You can configure self-signed service certificates with your edge gateways, to use in their VPN-related capabilities. You can create, install, and manage self-signed certificates.

If the service certificate is available on the Certificates screen, you can specify that service certificate when you configure the VPN-related settings of the edge gateway. The VPN presents the specified service certificate to the clients accessing the VPN.

## Prerequisites

Verify that at least one CSR is available on the **Certificates** screen for the edge gateway. See [Generate a Certificate Signing Request for an Edge Gateway](#).

## Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Certificates** tab.
- 3 Select the CSR in the list that you want to use for this self-signed certificate and click **Self-sign CSR**.
- 4 Type the number of days that the self-signed certificate is valid for.
- 5 Click **Keep**.

The system generates the self-signed certificate and adds a new entry with type Service Certificate to the on-screen list.

## Results

The self-signed certificate is available on the edge gateway. In the on-screen list, when you select an entry with type Service Certificate, its details are displayed in the screen.

## Add a CA Certificate to the Edge Gateway for SSL Certificate Trust Verification

Adding a CA certificate to an edge gateway enables trust verification of SSL certificates that are presented to the edge gateway for authentication, typically the client certificates used in VPN connections to the edge gateway.

You usually add the root certificate of your company or organization as a CA certificate. A typical use is for SSL VPN, where you want to authenticate VPN clients using certificates. Client certificates can be distributed to the VPN clients and when the VPN clients connect, their client certificates are validated against the CA certificate.

---

**Note** When adding a CA certificate, you typically configure a relevant Certificate Revocation List (CRL). The CRL protects against clients that present revoked certificates. See [Add a Certificate Revocation List to an Edge Gateway](#).

---

### Prerequisites

Verify that you have the CA certificate data in PEM format. In the user interface, you can either paste in the PEM data of the CA certificate or browse to a file that contains the data and is available in your network from your local system.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Certificates** tab.
- 3 Click **CA certificate**.
- 4 Provide the CA certificate data.
  - If the data is in a PEM file on a system you can navigate to, click the **Upload** button to browse to the file and select it.
  - If you can copy and paste the PEM data, paste it into the **CA Certificate (PEM format)** field.
 

Include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` lines.
- 5 (Optional) Type a description.
- 6 Click **Keep**.

### Results

The CA certificate with type CA Certificate appears in the on-screen list. This CA certificate is now available for you to specify when you configure the VPN-related settings of the edge gateway.

## Add a Certificate Revocation List to an Edge Gateway

A Certificate Revocation List (CRL) is a list of digital certificates that the issuing Certificate Authority (CA) claims to be revoked, so that systems can be updated not to trust users that present those revoked certificates. You can add CRLs to the edge gateway.

As described in the *NSX Administration Guide*, the CRL contains the following items:

- The revoked certificates and the reasons for revocation

- The dates that the certificates are issued
- The entities that issued the certificates
- A proposed date for the next release

When a potential user attempts to access a server, the server allows or denies access based on the CRL entry for that particular user.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Certificates** tab.
- 3 Click **CRL**.
- 4 Provide the CRL data.
  - If the data is in a PEM file on a system you can navigate to, click the **Upload** button to browse to the file and select it.
  - If you can copy and paste the PEM data, paste it into the **CRL (PEM format)** field. Include the `-----BEGIN X509 CRL-----` and `-----END X509 CRL-----` lines.
- 5 (Optional) Type a description.
- 6 Click **Keep**.

#### Results

The CRL appears in the on-screen list.

## Add a Service Certificate to the Edge Gateway

Adding service certificates to an edge gateway makes those certificates available for use in the VPN-related settings of the edge gateway. You can add a service certificate to the **Certificates** screen.

#### Prerequisites

Verify that you have the service certificate and its private key in PEM format. In the user interface, you can either paste in the PEM data or browse to a file that contains the data and is available in your network from your local system.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.

- 2 Click the **Certificates** tab.
- 3 Click **Service certificate**.
- 4 Input the PEM-formatted data of the service certificate.
  - If the data is in a PEM file on a system you can navigate to, click the **Upload** button to browse to the file and select it.
  - If you can copy and paste the PEM data, paste it into the **Service Certificate (PEM format)** field.  
 Include the `-----BEGIN CERTIFICATE-----` and `-----END CERTIFICATE-----` lines.
- 5 Input the PEM-formatted data of the certificate private key.  
 When FIPS mode is on, RSA key sizes must be greater or equal to 2048 bits.
  - If the data is in a PEM file on a system you can navigate to, click the **Upload** button to browse to the file and select it.
  - If you can copy and paste the PEM data, paste it into the **Private Key (PEM format)** field.  
 Include the `-----BEGIN RSA PRIVATE KEY-----` and `-----END RSA PRIVATE KEY-----` lines.
- 6 Enter a private key passphrase and confirm it.
- 7 (Optional) Enter a description.
- 8 Click **Keep**.

### Results

The certificate with type Service Certificate appears in the on-screen list. This service certificate is now available for you to select when you configure the VPN-related settings of the edge gateway.

## Custom Grouping Objects for NSX Data Center for vSphere Edge Gateways

The NSX Data Center for vSphere software in your VMware Cloud Director environment provides the capability for defining sets and groups of certain entities, which you can then use when specifying other network-related configurations, such as in firewall rules.

### Create an IP Set for Use in Firewall Rules and DHCP Relay Configuration

An IP set is a group of IP addresses that you can create at an organization virtual data center level. You can use an IP set as the source or destination in a firewall rule or in a DHCP relay configuration.

You create an IP set by using the **Grouping Objects** page of the VMware Cloud Director tenant portal. The **Grouping Objects** page is available on both the Services and Edge Gateway screens.

## Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
Open through Edge Gateway Services	<ol style="list-style-type: none"> <li>a Navigate to <b>Networking &gt; Edges</b>.</li> <li>b Select the edge gateway that you want to edit, and click <b>Configure Services</b>.</li> <li>c Click <b>Grouping Objects</b>.</li> </ol>
Open through Security Services	<ol style="list-style-type: none"> <li>a Navigate to <b>Networking &gt; Security</b>.</li> <li>b Select the security service that you want to edit, and click <b>Configure Services</b>.</li> <li>c Click <b>Grouping Objects</b>.</li> </ol>

- 2 Click the **IP Sets** tab.

The IP sets that are already defined are displayed on the screen.

- 3 To add an IP set, click the **Create** () button.
- 4 Enter a name, optionally, a description for the IP set, and the IP addresses to be included in the set.
- 5 (Optional) If you are specifying the IP set using the **Grouping Objects** page on the Services screen, use the **Inheritance** toggle to enable inheritance and allow visibility at the underlying scopes.

Inheritance is enabled by default.

- 6 To save this IP set, click **Keep**.

## Results

The new IP set is available for selection as the source or destination in firewall rules or in DHCP relay configurations.

## Create a MAC Set for Use in Firewall Rules

An MAC set is a group of MAC addresses that you can create at an organization virtual data center level. You can use a MAC set as the source or destination in a firewall rule.

You create a MAC set using the **Grouping Objects** page of the VMware Cloud Director tenant portal. The Grouping Objects page is available on both the **Services** and **Edge Gateway** screens.

## Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
Open through Edge Gateway Services	<ol style="list-style-type: none"> <li>Navigate to <b>Networking &gt; Edges</b>.</li> <li>Select the edge gateway that you want to edit, and click <b>Configure Services</b>.</li> <li>Click <b>Grouping Objects</b>.</li> </ol>
Open through Security Services	<ol style="list-style-type: none"> <li>Navigate to <b>Networking &gt; Security</b>.</li> <li>Select the security service that you want to edit, and click <b>Configure Services</b>.</li> <li>Click <b>Grouping Objects</b>.</li> </ol>

- 2 Click the **MAC Sets** tab.

The MAC sets that are already defined are displayed on the screen.

- 3 To add a MAC set, click the **Create** () button.
- 4 Enter a name for the set, optionally, a description, and the MAC addresses to be included in the set.
- 5 (Optional) If you are specifying the MAC set using the **Grouping Objects** page on the **Services** screen, use the **Inheritance** toggle to enable inheritance and allow visibility at underlying scopes.

Inheritance is enabled by default.

- 6 To save the MAC set, click **Keep**.

## Results

The new MAC set is available for selection as the source or destination in firewall rules.

## View Services Available for Firewall Rules

You can view the list of services that are available for use in firewall rules. In this context, a service is a protocol-port combination.

You can view the available services using the Grouping Objects page of the VMware Cloud Director tenant portal. The Grouping Objects page is available on both the Services and Edge Gateway screens.

You cannot add new services to the list using the tenant portal. The set of services available for your use is managed by your VMware Cloud Director system administrator.

## Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
Open through Edge Gateway Services	<ol style="list-style-type: none"> <li>Navigate to <b>Networking &gt; Edges</b>.</li> <li>Select the edge gateway that you want to edit, and click <b>Configure Services</b>.</li> <li>Click <b>Grouping Objects</b>.</li> </ol>
Open through Security Services	<ol style="list-style-type: none"> <li>Navigate to <b>Networking &gt; Security</b>.</li> <li>Select the security service that you want to edit, and click <b>Configure Services</b>.</li> <li>Click <b>Grouping Objects</b>.</li> </ol>

- 2 Click the **Services** tab.

## Results

The available services are displayed on the screen.

## View Service Groups Available for Firewall Rules

You can view the list of service groups that are available for use in firewall rules. In this context, a service is a protocol-port combination, and a service group is a group of services or other service groups.

You can view the available service groups using the Grouping Objects page of the VMware Cloud Director tenant portal. The Grouping Objects page is available on both the Services and Edge Gateway screens.

You cannot create service groups using the tenant portal. The set of service groups available for your use is managed by your VMware Cloud Director system administrator.

## Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
Open through Edge Gateway Services	<ol style="list-style-type: none"> <li>Navigate to <b>Networking &gt; Edges</b>.</li> <li>Select the edge gateway that you want to edit, and click <b>Configure Services</b>.</li> <li>Click <b>Grouping Objects</b>.</li> </ol>
Open through Security Services	<ol style="list-style-type: none"> <li>Navigate to <b>Networking &gt; Security</b>.</li> <li>Select the security service that you want to edit, and click <b>Configure Services</b>.</li> <li>Click <b>Grouping Objects</b>.</li> </ol>

- 2 Click the **Service Groups** tab.

## Results

The available service groups are displayed on the screen. The Description column displays the services that are grouped in each service group.

## Statistics and Logs for an NSX Data Center for vSphere Edge Gateway

You can view statistics and logs for an NSX Data Center for vSphere edge gateway.

### View Statistics

You can view statistics on the **Edge Gateway Services** screen.

#### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Statistics** tab.
- 3 Navigate through the tabs depending on the type of statistics you want to see.

Option	Description
<b>Connections</b>	The Connections screen provides operational visibility. The screen displays graphs for the traffic flowing through the interfaces of the selected edge gateway and for the firewall. Select the period for which you want to view the statistics.
<b>IPsec VPN</b>	The IPsec VPN screen displays the IPsec VPN status and statistics, and status and statistics for each tunnel.
<b>L2 VPN</b>	The L2 VPN screen displays the L2 VPN status and statistics.

### Enable Logging

You can enable logging for an edge gateway. In addition to enabling logging for the features for which you want to collect log data, to complete the configuration, you must have a Syslog server to receive the collected log data. When you configure a Syslog server on the Edge Settings screen, you are able to access the logged data from that Syslog server.

#### Prerequisites

- Verify that you are an **organization administrator** or you are assigned a role that includes an equivalent set of rights.
- Verify that your role includes the **Configure System Logging** right.

## Procedure

### 1 Open Edge Gateway Services.

- a In the top navigation bar, click **Networking** and click **Edge Gateways**.
- b Select the edge gateway that you want to edit and click **Services**.

### 2 On the **Edge Settings** tab, click the **Edit Syslog server** button.

You can customize the Syslog server for the networking-related logs of your edge gateway for those services that have logging enabled.

If the VMware Cloud Director system administrator has already configured a Syslog server for the VMware Cloud Director environment, the system uses that Syslog server by default and its IP address is displayed on the **Edge Settings** screen.

### 3 Enable logging per feature.

- On the **NAT** tab, click the **DNAT Rule** button, and turn on the **Enable logging** toggle.

Logs the address translation.

- On the **NAT** tab, click the **SNAT Rule** button, and turn on the **Enable logging** toggle.

Logs the address translation.

- On the **Routing** tab, click **Routing Configuration**, and under Dynamic Routing Configuration, turn on the **Enable logging** toggle.

Logs the dynamic routing activities. From the **Log Level** drop-down menu, you can select the lower bound of the message status level to log.

- On the **Load Balancer** tab, click **Global Configuration**, and turn on the **Enable logging** toggle.

Logs the traffic flow for the load balancer. From the **Log Level** drop-down menu, you can select the lower bound of the message status level to log.

- On the **VPN** tab, navigate to **IPSec VPN > Logging Settings**, and turn on the **Enable logging** toggle.

Logs the traffic flow between the local subnet and peer subnet. From the **Log Level** drop-down menu, you can select the lower bound of the message status level to log.

- On the **SSL VPN-Plus** tab, click **General Settings**, and turn on the **Enable logging** toggle.

Maintains a log of the traffic passing through the SSL VPN gateway.

- On the **SSL VPN-Plus** tab, click **Server Settings**, and turn on the **Enable logging** toggle.

Logs the activities that occur on the SSL VPN server, for Syslog. From the **Log Level** drop-down menu, you can select the lower bound of the message status level to log.

## Enable SSH Command-Line Access to an NSX Data Center for vSphere Edge Gateway

You can enable SSH command-line access to an edge gateway.

### Procedure

- 1 Open Edge Gateway Services.
  - a In the top navigation bar, click **Networking** and click **Edge Gateways**.
  - b Select the edge gateway that you want to edit and click **Services**.
- 2 Click the **Edge Settings** tab.
- 3 Configure the SSH settings.

Option	Description
<b>Username</b>	Enter the credentials for the SSH access to this edge gateway.
<b>Password</b>	By default, the SSH user name is <b>admin</b> .
<b>Retype Password</b>	
<b>Password Expiry</b>	Enter the expiration period for the password, in days.
<b>Login Banner</b>	Enter the text to be displayed to users when they begin an SSH connection to the edge gateway.

- 4 Turn on the **Enabled** toggle.

### What to do next

Configure the appropriate NAT or firewall rules to allow an SSH access to this edge gateway.

## Working with Security Tags for NSX Data Center for vSphere Edge Gateways

Security tags are labels which can be associated with a virtual machine or a group of virtual machines. Security tags are designed to be used with security groups. Once you create the security tags, you associate them with a security group which can be used in firewall rules. You can create, edit, or assign a user-defined security tag. You can also view which virtual machines or security groups have a particular security tag applied.

A common use case for security tags is to dynamically group objects to simplify firewall rules. For example, you might create several different security tags based on the type of activity you expect to occur on a given virtual machine. You create a security tag for database servers and another one for email servers. Then you apply the appropriate tag to virtual machines that house database servers or email servers. Later, you can assign the tag to a security group, and write a firewall rule against it, applying different security settings depending on whether the virtual machine is running a database server or an email server. Later, if you change the functionality of the virtual machine, you can remove the virtual machine from the security tag rather than editing the firewall rule.

## Create and Assign Security Tags

You can create a security tag and assign it to a virtual machine or a group of virtual machines.

You create a security tag and assign it to a virtual machine or a group of virtual machines.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.
- 2 Select a security service and click **Configure services**.
- 3 Click the **Security Tags** tab.

4 Click the **Create** () button, and enter a name for the security tag.

5 (Optional) Enter a description for the security tag.

6 (Optional) Assign the security tag to a virtual machine or a group of virtual machines.

In the **Browse objects of type** drop-down menu, **Virtual Machines** is selected by default.

- a Select a virtual machine from the left panel.
- b Assign the security tag to the selected virtual machine by clicking the right arrow.

The virtual machine moves to the right panel and is assigned the security tag.

7 When you complete assigning the tag to the selected virtual machines, click **Keep**.

### Results

The security tag is created, and if you chose, is assigned to selected virtual machines.

### What to do next

Security tags are designed to work with a security group. For more information about creating security groups, see [Create a Security Group](#).

## Change the Security Tag Assignment

After you create a security tag, you can manually assign it to virtual machines. You can also edit a security tag to remove the tag from the virtual machines to which you have already assigned it.

If you have created security tags, you can assign them to virtual machines. You can use security tags to group virtual machines for writing firewall rules. For example, you might assign a security tag to a group of virtual machines with highly sensitive data.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.
- 2 Select a security service and click **Configure services**.
- 3 Click the **Security Tags** tab.

- 4 From the list of security tags, select the security tag that you want to edit, and click the **Edit** button.
- 5 Select virtual machines from the left panel, and assign the security tag to them by clicking the right arrow.

The virtual machines in the right panel are assigned the security tag.

- 6 Select virtual machines in the right panel, and remove the tag from them by clicking the left arrow.

The virtual machines in the left panel do not have the security tag assigned.

- 7 When you finish adding your changes, click **Keep**.

### Results

The security tag is assigned to the selected virtual machines.

### What to do next

Security tags are designed to work with a security group. For more information about creating security groups, see [Create a Security Group](#).

## View Applied Security Tags

You can view the security tags applied to virtual machines in your environment. You can also see the security tags that are applied to security groups in your environment.

### Prerequisites

A security tag must have been created and applied to a virtual machine or to a security group.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.
- 2 Select a security service and click **Configure services**.
- 3 View the assigned tags from the **Security Tags** tab.
  - a On the **Security Tags** tab, select the security tag for which you want to see assignments, and click the **Edit** icon.
  - b Under the **Assign/Unassign VMs**, you can see the list of virtual machines assigned to the security tag.
  - c Click **Discard**.

- 4 View the assigned tags from the **Security Groups** tab.
  - a Click the **Grouping Objects** tab, and click **Security Groups**.
  - b Select a security group.
  - c From the list under **Include Members**, you can see the security tag assigned to a security group.

### Results

You can view the existing security tags and associated virtual machines and security groups. This way, you can determine a strategy for creating firewall rules based on security tags and security groups.

## Edit a Security Tag

You can edit a user-defined security tag.

If you change the environment or function of a virtual machine, you might also want to use a different security tag so that firewall rules are correct for the new machine configuration. For example, if you have a virtual machine where you no longer store sensitive data, you might want to assign a different security tag so that firewall rules that apply to sensitive data is no longer run against the virtual machine.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.
- 2 Select a security service and click **Configure services**.
- 3 Click the **Security Tags** tab.
- 4 From the list of security tags, select the security tag that you want to edit.
- 5 Click the **Edit** button.
- 6 Edit the name and the description of the security tag.
- 7 Assign the tag to or remove the assignment from the virtual machines that you select.
- 8 To save your changes, click **Keep**.

### What to do next

If you edit a security tag, you might also need to edit an associated security group or firewall rules. For more information about security groups, see [Working with Security Groups for NSX Data Center for vSphere Edge Gateways](#)

.

## Delete a Security Tag

You can delete a user-defined security tag.

You might want to delete a security tag if the function or environment of the virtual machine changes. For example, if you have a security tag for Oracle databases, but you decide to use a different database server, you can remove the security tag so that firewall rules that apply to Oracle databases no longer run against the virtual machine.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Networking**, select **Security**.
- 2 Select a security service and click **Configure services**.
- 3 Click the **Security Tags** tab.
- 4 From the list of security tags, select the security tag that you want to delete.
- 5 Click the **Delete** button.
- 6 To confirm the deletion, click **OK**.

#### Results

The security tag is deleted.

#### What to do next

If you delete a security tag, you might also need to edit an associated security group or firewall rules. For more information about security groups, see [Working with Security Groups for NSX Data Center for vSphere Edge Gateways](#).

## Working with Security Groups for NSX Data Center for vSphere Edge Gateways

A security group is a collection of assets or grouping objects, such as virtual machines, organization virtual data center networks, or security tags.

Security groups can have dynamic membership criteria based on security tags, virtual machine name, virtual machine guest OS name, or virtual machine guest host name. For example, all virtual machines that have the security tag "web" will be automatically added to a specific security group destined for Web servers. After creating a security group, a security policy is applied to that group.

### Create a Security Group

You can create user-defined security groups.

#### Prerequisites

If you want to use security tags with security groups, [Create and Assign Security Tags](#).

## Procedure

- 1 Open the Security Services.
  - a Navigate to **Networking > Security**.
  - b Select the organization VDC for which you want to apply security settings, and click **Configure Services**.

The tenant portal opens Security Services.

- 2 Navigate to **Grouping Objects > Security Groups**

The **Security Groups** page opens.

- 3 Click the **Create** () button.

- 4 Enter a name and, optionally, a description for the security group.

The description displays in the list of security groups, so adding a meaningful description can make it easy to identify the security group at a glance.

- 5 (Optional) Add a dynamic member set.

- a Click the **Add** () button under Dynamic Member Sets.

- b Select whether to match **Any** or **All** of the criteria in your statement.

- c Enter the first object to match.

The options are **Security Tag**, **VM Guest OS Name**, **VM Name**, and **VM Guest Host Name**.

- d Select an operator, such as **Contains**, **Starts with**, or **Ends with**.

- e Enter a value.

- f (Optional) To add another statement, use a Boolean operator **And** or **Or**.

- 6 (Optional) Include Members.

- a From the **Browse objects of type** drop-down menu, select the type of objects, such as **Virtual Machines**, **Org VDC networks**, **IP sets**, **MAC sets**, or **Security tags**.

- b To include an object in the Include Members list, select the object from the left panel, and move it to the right panel by clicking the right arrow.

- 7 (Optional) Exclude members.

- a From the **Browse objects of type** drop-down menu, select the type of objects, such as **Virtual Machines**, **Org VDC networks**, **IP sets**, **MAC sets**, or **Security tags**.

- b To include an object in the Exclude Members list, select the object from the left panel, and move it to the right panel by clicking the right arrow.

- 8 To preserve your changes, click **Keep**.

## Results

The security group can now be used in rules, such as firewall rules.

## Edit a Security Group

You can edit user-defined security groups.

### Procedure

1 Open the Security Services.

- a Navigate to **Networking > Security**.
- b Select the organization VDC for which you want to apply security settings, and click **Configure Services**.

The tenant portal opens Security Services.

2 Navigate to **Grouping Objects > Security Groups**

The **Security Groups** page opens.

3 Select the security group you want to edit.

The details for the security group display below the list of security groups.

4 (Optional) Edit the name and the description of the security group.

5 (Optional) Add a dynamic member set.

- a Click the **Add** button under **Dynamic Member Sets**.
- b Select whether to match **Any** or **All** of the criteria in your statement.
- c Enter the first object to match.

The options are **Security Tag**, **VM Guest OS Name**, **VM Name**, and **VM Guest Host Name**.

d Select an operator, such as **Contains**, **Starts with**, or **Ends with**.

e Enter a value.

f (Optional) To add another statement, use a Boolean operator **And** or **Or**.

6 (Optional) Edit a dynamic member set by clicking the **Edit** icon next to the member set that you want to edit.

- a Apply the necessary changes to the dynamic member set.
- b Click **OK**.

7 (Optional) Delete a dynamic member set by clicking the **Delete** icon next to the member set that you want to delete.

- 8 (Optional) Edit the included members list by clicking the **Edit** icon next to the Include Members list.
  - a From the **Browse objects of type** drop-down menu, select the type of objects, such as **Virtual Machines, Org VDC networks, IP sets, MAC sets, or Security tags**.
  - b To include an object in the include members list, select the object from the left panel, and move it to the right panel by clicking the right arrow.
  - c To exclude an object from the include members list, select the object from the right panel, and move it to the left panel by clicking the left arrow.
- 9 (Optional) Edit the excluded members list by clicking the **Edit** icon next to the Exclude Members list.
  - a From the **Browse objects of type** drop-down menu, select the type of objects, such as **Virtual Machines, Org VDC networks, IP sets, MAC sets, or Security tags**.
  - b To include an object in the exclude members list, select the object from the left panel, and move it to the right panel by clicking the right arrow.
  - c To exclude an object from the exclude members list, select the object from the right panel, and move it to the left panel by clicking the left arrow.
- 10 Click **Save changes**.
 

The changes to the security group are saved.

## Delete a Security Group

You can delete a user-defined security group.

### Procedure

- 1 Open the Security Services.
  - a Navigate to **Networking > Security**.
  - b Select the organization VDC for which you want to apply security settings, and click **Configure Services**.
 

The tenant portal opens Security Services.
- 2 Navigate to **Grouping Objects > Security Groups**

The **Security Groups** page opens.
- 3 Select the security group you want to delete.
- 4 Click the **Delete** button.
- 5 To confirm the deletion, click **OK**.

### Results

The security group is deleted.

## Managing NSX-T Data Center Edge Gateways

An NSX-T Data Center edge gateway provides a routed organization VDC network or a data center group network with connectivity to external networks and IP management properties. It can also provide services such as firewall, network address translation (NAT), IPSec VPN, DNS forwarding, and DHCP, which is enabled by default.

### Dedicated External Networks

To provide a fully routed network topology in a virtual data center, your **system administrator** can dedicate an external network to a specific NSX-T Data Center edge gateway.

In this configuration, there is a one-to-one relationship between the external network and the NSX-T Data Center edge gateway, and other edge gateways cannot connect to the external network.

An NSX-T Data Center tier-0 logical router or VRF gateway that is associated with a dedicated external network is part of the tenant networking stack. The external network is considered a part of the VMware Cloud Director network routing domain.

A dedicated external network provides additional edge gateway routing services, such as route advertisement management and border gateway protocol (BGP) configuration.

You can decide which of the networks that are attached to the edge gateway to advertise to the external network. This makes possible a mixture of NAT-routed and fully routed organization virtual data center networks.

### Add an IP Set to an NSX-T Data Center Edge Gateway

To create firewall rules and add them to an NSX-T Data Center edge gateway, you must first create IP sets. IP sets are groups of objects to which the firewall rules apply. Combining multiple objects into IP sets helps reduce the total number of firewall rules to be created.

#### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the NSX-T edge gateway.
- 3 Under **Security**, click **IP Sets** tab and click **New**.
- 4 Enter a name and, optionally, a description for the IP set.
- 5 Enter an IP address or an IP addresses range for the virtual machines that the IP set includes, and click **Add**.
- 6 To save the firewall group, click **Save**.

#### Results

You created an IP set and added it to the NSX-T edge gateway.

**What to do next**[Add an NSX-T Data Center Edge Gateway Firewall Rule](#)

## Add an NSX-T Data Center Edge Gateway Firewall Rule

To control the incoming and outgoing network traffic to and from an NSX-T Data Center edge gateway, you create firewall rules.

**Procedure**

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 If the **Firewall** screen is not already visible under the Services section, click the **Firewall** tab.
- 4 Click **Edit Rules**.
- 5 Click the **New On Top** button.

A row for the new rule is added above the selected rule.

- 6 Configure the firewall rule.

Option	Description
<b>Name</b>	Enter a name for the rule.
<b>State</b>	To enable the rule upon creation, turn on the <b>State</b> toggle.
<b>Applications</b>	(Optional) To select a specific port profile to which the rule applies, turn on the <b>Applications</b> toggle and click <b>Save</b> .
<b>Source</b>	Select an option and click <b>Keep</b> . <ul style="list-style-type: none"> <li>■ To allow or deny traffic from any source address, toggle on <b>Any Source</b>.</li> <li>■ To allow or deny traffic from specific firewall groups, select the firewall groups from the list.</li> </ul>
<b>Destination</b>	Select an option and click <b>Keep</b> . <ul style="list-style-type: none"> <li>■ To allow or deny traffic to any destination address, toggle on <b>Any Destination</b>.</li> <li>■ To allow or deny traffic to specific firewall groups, select the firewall groups from the list.</li> </ul>
<b>Action</b>	From the <b>Action</b> drop-down menu, select an option. <ul style="list-style-type: none"> <li>■ To allow traffic from or to the specified sources, destinations, and services, select <b>Accept</b>.</li> <li>■ To block traffic from or to the specified sources, destinations, and services, without notifying the blocked client select <b>Drop</b>.</li> <li>■ To block traffic from or to the specified sources, destinations, and services, and to notify the blocked client that traffic was rejected, select <b>Reject</b>.</li> </ul>
<b>IP Protocol</b>	Select whether to apply the rule to IPv4 or IPv6 traffic.

Option	Description
Direction	Select the traffic direction to which to apply the rule.  <b>Note</b> In VMware Cloud Director 10.2.1 and later versions, this option is no longer available.
Enable logging.	To have the address translation performed by this rule logged, turn on the <b>Enable logging</b> toggle.

- 7 Click **Save**.
- 8 To configure additional rules, repeat these steps.

### Results

After the firewall rules are created, they appear in the Edge Gateway Firewall Rules list. You can move up, move down, edit, or delete the rules as needed.

## Add an SNAT or a DNAT Rule to an NSX-T Edge Gateway

To change the source IP address from a private to a public IP address, you create a source NAT (SNAT) rule. To change the destination IP address from a public to a private IP address, you create a destination NAT (DNAT) rule.

When you configure a SNAT or a DNAT rule on an edge gateway in the VMware Cloud Director environment, you always configure the rule from the perspective of your organization VDC.

An SNAT rule translates the source IP address of packets sent from an organization VDC network out to an external network or to another organization VDC network.

A NO SNAT rule prevents the translation of the internal IP address of packets sent from an organization VDC out to an external network or to another organization VDC network.

A DNAT rule translates the IP address and, optionally, the port of packets received by an organization VDC network that are coming from an external network or from another organization VDC network.

A NO DNAT rule prevents the translation of the external IP address of packets received by an organization VDC from an external network or from another organization VDC network.

VMware Cloud Director supports automatic route redistribution when you use NAT services on an NSX-T Data Center edge gateway.

---

**Important** If you are using Tanzu Kubernetes clusters, make note of the system SNAT rule created on the edge gateway to avoid creating a conflicting rule.

---

### Prerequisites

The public IP addresses must have been added to the edge gateway interface on which you want to add the rule.

## Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway and, under **Services**, click **NAT**.
- 3 To add a rule, click **New**.
- 4 Configure an SNAT or NO SNAT rule (inside going outside).

Option	Description
<b>Name</b>	Enter a meaningful name for the rule.
<b>Description</b>	(Optional) Enter a description for the rule.
<b>Interface type</b>	From the drop-down menu, select SNAT or NO SNAT.
<b>External IP</b>	Depending on the type of rule that you are creating, choose one of the options. <ul style="list-style-type: none"> <li>■ If you are creating a SNAT rule, enter the public IP address of the edge gateway for which you are configuring the SNAT rule.</li> <li>■ If you are creating a NO SNAT rule, leave the text box empty.</li> </ul>
<b>Internal IP</b>	Enter the IP address or a list of IP addresses of the virtual machines for which you are configuring SNAT, so that they can send traffic to the external network.
<b>Destination IP</b>	(Optional) If you want the rule to apply only for traffic to a specific domain, enter an IP address for this domain or an IP address list. If you leave this text box blank, the SNAT rule applies to all destinations outside of the local subnet.
<b>Advanced Settings (Optional)</b>	Click the <b>Advanced Settings</b> tab for some additional settings. <p><b>State</b></p> <p>To activate the rule upon creation, toggle on the <b>State</b> option.</p> <p><b>Logging</b></p> <p>To have the address translation performed by this rule logged, toggle on the <b>Logging</b> option.</p> <p><b>Priority</b></p> <p>If an address has multiple NAT rules, you can assign these rules different priorities to determine the order in which they are applied. A lower value means a higher priority for this rule.</p> <p><b>Firewall Match</b></p> <p>You can set a firewall match rule to determine how firewall is applied during NAT. From the drop-down menu, select one of the following options.</p> <ul style="list-style-type: none"> <li>■ To apply firewall rules to the internal address of a NAT rule, select <b>Match Internal Address</b>.</li> <li>■ To apply firewall rules to the external address of a NAT rule, select <b>Match External Address</b>.</li> <li>■ To skip applying firewall rules, select <b>Bypass</b>.</li> </ul>

## 5 Configure a DNAT or NO DNAT rule (outside going inside).

Option	Description
<b>Name</b>	Enter a meaningful name for the rule.
<b>Description</b>	(Optional) Enter a description for the rule.
<b>Interface type</b>	From the drop-down menu, select DNAT or NO DNAT.
<b>External IP</b>	Enter the public IP address of the edge gateway for which you are configuring the DNAT rule. The IP addresses that you enter must be suballocated to the edge gateway.
<b>External Port</b>	(Optional) Enter a port into which the DNAT rule is translating for the packets inbound to the virtual machines.
<b>Internal IP</b>	Depending on the type of rule that you are creating, choose one of the options. <ul style="list-style-type: none"> <li>■ If you are creating a DNAT rule, enter the IP address or a list of IP address of the virtual machines for which you are configuring DNAT, so that they can receive traffic from the external network.</li> <li>■ If you are creating a NO DNAT rule, leave the text box empty.</li> </ul>
<b>Application</b>	(Optional) Select a specific application port profile to which to apply the rule. The application port profile includes a port and a protocol that the incoming traffic uses on the edge gateway to connect to the internal network.
<b>Advanced Settings (Optional)</b>	Click the <b>Advanced Settings</b> tab for some additional settings. <p><b>State</b></p> <p>To activate the rule upon creation, toggle on the <b>State</b> option.</p> <p><b>Logging</b></p> <p>To have the address translation performed by this rule logged, toggle on the <b>Logging</b> option.</p> <p><b>Priority</b></p> <p>If an address has multiple NAT rules, you can assign these rules different priorities to determine the order in which they are applied. A lower value means a higher priority for this rule.</p> <p><b>Firewall Match</b></p> <p>You can set a firewall match rule to determine how firewall is applied during NAT. From the drop-down menu, select one of the following options.</p> <ul style="list-style-type: none"> <li>■ To apply firewall rules to the internal address of a NAT rule, select <b>Match Internal Address</b>.</li> <li>■ To apply firewall rules to the external address of a NAT rule, select <b>Match External Address</b>.</li> <li>■ To skip applying firewall rules, select <b>Bypass</b>.</li> </ul>

## 6 Click **Save**.

## 7 To configure additional rules, repeat these steps.

## Configure a DNS Forwarder Service on an NSX-T Edge Gateway

To forward DNS queries to external DNS servers, configure a DNS forwarder.

As part of your DNS forwarder service configuration, you can also add conditional forwarder zones. A conditional forwarder zone is configured as a list containing up to five FQDN DNS zones. If a DNS query matches a domain name from that list, the query is forwarded to the servers from the corresponding forwarder zone.

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway and, under **IP Management**, click **DNS**.
- 3 In the **DNS Forwarder** section, click **Edit**.
- 4 To enable the DNS Forwarder service, turn on the **State** toggle.
- 5 Enter a name and, optionally, a description for the default DNS zone.
- 6 Enter one or more upstream server IP addresses, separated by a comma.
- 7 Click **Save**.
- 8 (Optional) Add a conditional forwarder zone.
  - a In the **Conditional Forwarder Zone** section, click **Add**.
  - b Enter a name for the forwarder zone.
  - c Enter one or more upstream server IP addresses, separated by a comma.
  - d Enter one or more domain names, separated by a comma, and click **Save**.

## Create Custom Application Port Profiles

To create firewall and NAT rules, you can use preconfigured application port profiles and custom application port profiles.

Application port profiles include a combination of a protocol and a port, or a group of ports, that is used for firewall and NAT services on the edge gateway. In addition to the default port profiles that are preconfigured for NSX-T Data Center, you can create custom application port profiles.

When you create a custom application port profile on an edge gateway, it becomes visible to all the other NSX-T Data Center edge gateways that are in the same organization VDC.

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Security**, click **Application Port Profiles**.
- 4 In the **Custom Applications** section, click **New**.
- 5 Enter a name and, optionally, a description for the application port profile.

- 6 Select a protocol from the drop-down menu.
- 7 Enter a port, or a range of ports, separated by a comma, and click **Save**.

#### What to do next

Use application port profiles to create firewall and NAT rules. See [Add an NSX-T Data Center Edge Gateway Firewall Rule](#) and [Add an SNAT or a DNAT Rule to an NSX-T Edge Gateway](#).

## IPsec Policy-Based VPN for NSX-T Data Center Edge Gateways

Starting with version 10.1, VMware Cloud Director supports site-to-site policy-based IPsec VPN between an NSX-T Data Center edge gateway instance and a remote site.

IPsec VPN offers site-to-site connectivity between an edge gateway and remote sites which also use NSX-T Data Center or which have either third-party hardware routers or VPN gateways that support IPsec.

Policy-based IPsec VPN requires a VPN policy to be applied to packets to determine which traffic is to be protected by IPsec before being passed through a VPN tunnel. This type of VPN is considered static because when a local network topology and configuration change, the VPN policy settings must also be updated to accommodate the changes.

NSX-T Data Center edge gateways support split tunnel configuration, with IPsec traffic taking routing precedence.

VMware Cloud Director supports automatic route redistribution when you use IPsec VPN on an NSX-T edge gateway.

### Configure NSX-T Policy-Based IPsec VPN

You can configure site-to-site connectivity between an NSX-T Data Center edge gateway and remote sites. The remote sites must use NSX-T Data Center, have third-party hardware routers, or VPN gateways that support IPsec.

VMware Cloud Director supports automatic route redistribution when you configure IPsec VPN on an NSX-T Data Center edge gateway.

#### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Services**, click **IPsec VPN**.
- 4 To configure an IPsec VPN tunnel, click **New**.
- 5 Enter a name and, optionally, a description for the IPsec VPN tunnel.
- 6 To enable the tunnel upon creation, toggle on the **Enabled** option.
- 7 Choose a pre-shared key to enter.

---

**Note** The pre-shared key must be the same on the other end of the IPsec VPN tunnel.

---

- 8 Enter one of the IP addresses that are available to the edge gateway for the local endpoint.

---

**Note** The IP address must be either the primary IP of the edge gateway, or an IP address that is separately allocated to the edge gateway from the external network .

---

- 9 Enter at least one local IP subnet address in CIDR notation to use for the IPsec VPN tunnel.
- 10 Enter the IP address for the remote site.
- 11 Enter at least one remote IP subnet address in CIDR notation to use for the IPsec VPN tunnel.
- 12 (Optional) To enable logging, toggle on the **Logging** option.
- 13 Click **Save**.
- 14 To verify that the tunnel is functioning, select it and click **View Statistics**.

If the tunnel is functioning, **Tunnel Status** and **IKE Service Status** both display **Up**.

### Results

The newly created IPsec VPN tunnel is listed in the **IPsec VPN** view. The IPsec VPN tunnel is created with a default security profile.

### What to do next

You can edit the IPsec VPN tunnel settings and customize its security profile as needed.

## Customize the Security Profile of an IPsec VPN Tunnel

If you decide not to use the system-generated security profile that was assigned to your IPsec VPN tunnel upon creation, you can customize it.

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Services**, click **IPsec VPN**.
- 4 Select the IPsec VPN tunnel and click **Security Profile Customization**.

## 5 Configure the IKE profiles.

The Internet Key Exchange (IKE) profiles provide information about the algorithms that are used to authenticate, encrypt, and establish a shared secret between network sites when you establish an IKE tunnel.

- a Select an IKE protocol version to set up a security association (SA) in the IPSec protocol suite.

Option	Description
<b>IKEv1</b>	When you select this option, IPSec VPN initiates and responds to IKEv1 protocol only.
<b>IKEv2</b>	The default option. When you select this version, IPSec VPN initiates and responds to IKEv2 protocol only.
<b>IKE-Flex</b>	When you select this option, if the tunnel establishment fails with IKEv2 protocol, the source site does not fall back and initiate a connection with the IKEv1 protocol. Instead, if the remote site initiates a connection with the IKEv1 protocol, then the connection is accepted.

- b Select a supported encryption algorithm to use during the Internet Key Exchange (IKE) negotiation.
  - c From the **Digest** drop-down menu, select a secure hashing algorithm to use during the IKE negotiation.
  - d From the **Diffie-Hellman Group** drop-down menu, select one of the cryptography schemes that allows the peer site and the edge gateway to establish a shared secret over an insecure communications channel.
  - e (Optional) In the **Association Lifetime** text box, modify the default number of seconds before the IPSec tunnel needs to reestablish.
- ## 6 Configure the IPSec VPN tunnel.

- a To enable perfect forward secrecy, toggle on the option.
- b Select a defragmentation policy.

The defragmentation policy helps to handle defragmentation bits present in the inner packet.

Option	Description
<b>Copy</b>	Copies the defragmentation bit from the inner IP packet to the outer packet.
<b>Clear</b>	Ignores the defragmentation bit present in the inner packet.

- c Select a supported encryption algorithm to use during the Internet Key Exchange (IKE) negotiation.
- d From the **Digest** drop-down menu, select a secure hashing algorithm to use during the IKE negotiation.

- e From the **Diffie-Hellman Group** drop-down menu, select one of the cryptography schemes that allows the peer site and the edge gateway to establish a shared secret over an insecure communications channel.
  - f (Optional) In the **Association Lifetime** text box, modify the default number of seconds before the IPSec tunnel needs to reestablish.
- 7 (Optional) In the **Probe Interval** text box, modify the default number of seconds for dead peer detection.
  - 8 Click **Save**.

## Results

In the IPSec VPN view, the security profile of the IPSec VPN tunnel displays as **User Defined**.

## Configure Dedicated External Network Services

To provide a fully routed network topology in a virtual data center, a **system administrator** can dedicate an external network to a specific NSX-T Data Center edge gateway.

When you use a dedicated external network, you can configure additional routing services, such as route advertisement management and border gateway protocol (BGP) configuration.

## Procedure

### 1 [Manage Route Advertisement](#)

By using route advertisement, you can create a fully routed network environment in an organization virtual data center (VDC).

### 2 [Configure BGP General Settings](#)

You can configure an external or internal Border Gateway Protocol (eBGP or iBGP) connection between an NSX-T Data Center edge gateway that has a dedicated external network and a router in your physical infrastructure.

### 3 [Create an IP Prefix List](#)

You can create IP prefix lists which contain single or multiple IP addresses. You use IP prefix lists to assign BGP neighbors with access permissions for route advertisement.

### 4 [Add a BGP Neighbor](#)

You can configure individual settings for the BGP routing neighbors when you add them.

## Manage Route Advertisement

By using route advertisement, you can create a fully routed network environment in an organization virtual data center (VDC).

You can decide which of the network subnets that are attached to the NSX-T Data Center edge gateway to advertise to the dedicated external network.

If a subnet is not added to the advertisement filter, the route to it is not advertised to the external network and the subnet remains private.

---

**Note** VMware Cloud Director advertises any organization VDC network that falls within the advertised route. Because of that, you do not need to create a filter for each subnet that is part of an advertised network.

---

Route advertisement is automatically configured on the NSX-T Data Center edge gateway.

VMware Cloud Director supports automatic route redistribution when you use route advertisement on an NSX-T edge gateway. Route redistribution is automatically configured on the tier-0 logical router which represents the dedicated external network.

### Prerequisites

- Verify that the **system administrator** dedicated an external network to an NSX-T Data Center edge gateway in your organization.
- Verify that you are an **organization administrator** or you are assigned a role that includes an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Routing**, click **Route Advertisement** and **Edit**.
- 4 To add a subnet to be advertised, click **Add**.
- 5 Add an IPv4 or IPv6 subnet.

Use the format *network\_gateway\_IP\_address/subnet\_prefix\_length*, for example, **192.167.1.1/24**.

## Configure BGP General Settings

You can configure an external or internal Border Gateway Protocol (eBGP or iBGP) connection between an NSX-T Data Center edge gateway that has a dedicated external network and a router in your physical infrastructure.

BGP makes core routing decisions by using a table of IP networks, or prefixes, which designate multiple routes between autonomous systems (AS).

The term BGP speaker refers to a networking device that is running BGP. Two BGP speakers establish a connection before any routing information is exchanged.

The term BGP neighbor refers to a BGP speaker that has established such a connection. After establishing the connection, the devices exchange routes and synchronize their tables. Each device sends keep-alive messages to keep this relationship alive.

---

**Note** In an edge gateway that is connected to an external network backed by a VRF gateway, the local AS number and graceful restart settings are read-only. Your **system administrator** can edit these settings on the parent tier-0 gateway in NSX-T Data Center.

---

### Prerequisites

- Verify that the **system administrator** dedicated an external network to an NSX-T Data Center edge gateway in your organization.
- Verify that you are an **organization administrator** or you are assigned a role that includes an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Routing**, click **BGP** and, under **Configuration**, click **Edit**.
- 4 Toggle on the **Status** option to activate BGP.
- 5 Enter an autonomous system (AS) ID number to use for the local AS feature of the protocol. VMware Cloud Director assigns the local AS number to the edge gateway. The edge gateway advertises this ID when it connects with its BGP neighbors in other autonomous systems.
- 6 From the drop-down menu, select a **Graceful Restart Mode** option.

Option	Description
<b>Helper and graceful restart</b>	<p>It is not a best practice to activate the graceful restart capability on the edge gateway because the BGP peerings from all gateways are always active. In case of a failover, the graceful restart capability increases the time a remote neighbor takes to select an alternate tier-0 gateway. This delays BFD-based convergence.</p> <hr/> <p><b>Note</b> The edge gateway configuration applies to all BGP neighbors unless the neighbor-specific configuration overrides it.</p>
<b>Helper only</b>	<p>Useful for reducing or eliminating the disruption of traffic associated with routes learned from a neighbor that is capable of graceful restart. The neighbor must be able to preserve its forwarding table while it undergoes a restart.</p>
<b>Disable</b>	<p>Deactivate graceful restart mode on the edge gateway.</p>

- 7 (Optional) Change the default value for the graceful restart timer.
- 8 (Optional) Change the default value for the stale route timer.
- 9 Toggle on the **ECMP** option to activate ECMP.

10 Click **Save**.

#### What to do next

- [Create an IP Prefix List](#)
- [Add a BGP Neighbor](#)

## Create an IP Prefix List

You can create IP prefix lists which contain single or multiple IP addresses. You use IP prefix lists to assign BGP neighbors with access permissions for route advertisement.

The IP prefix lists are referenced through BGP neighbor filters to limit the number of BGP updates that are exchanged between BGP peers. By using route filtering, you can reduce the amount of system resources needed for BGP updates.

For example, you can add the IP address 192.168.100.3/27 to the IP prefix list and deny the route from being redistributed to the edge gateway.

You can also append an IP address with `less than or equal to (le)` and `greater than or equal to (ge)` modifiers to grant or limit route redistribution. For example, `192.168.100.3/27 ge 26 le 32` modifiers match subnet masks greater than or equal to 26-bits and less than or equal to 32-bits in length.

#### Prerequisites

- Verify that the **system administrator** dedicated an external network to an NSX-T Data Center edge gateway in your organization.
- Verify that you are an **organization administrator** or you are assigned a role that includes an equivalent set of rights.
- [Configure BGP General Settings](#).

#### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Routing**, click **BGP** and **IP Prefix Lists**.
- 4 To add an IP prefix list, click **New**.
- 5 Enter a name and, optionally, a description for the prefix list.
- 6 Click **New** and add a CIDR notation for the prefix.
- 7 From the drop-down menu, select an action to apply to the prefix.
- 8 (Optional) Enter `greater than or equal to` and `less than or equal to` modifiers to grant or limit route redistribution.

## What to do next

- You can edit or delete the IP prefix list as needed.
- Configure route filtering. See [Add a BGP Neighbor](#).

## Add a BGP Neighbor

You can configure individual settings for the BGP routing neighbors when you add them.

### Prerequisites

- Verify that the **system administrator** dedicated an external network to an NSX-T Data Center edge gateway in your organization.
- Verify that you are an **organization administrator** or you are assigned a role that includes an equivalent set of rights.
- Verify that you configured the global BGP settings for the edge gateway. See [Configure BGP General Settings](#).
- If you use route filtering, verify that you created IP prefix lists. See [Create an IP Prefix List](#).

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the edge gateway.
- 3 Under **Routing**, click **BGP** and **Neighbors**.
- 4 To add a new BGP neighbor, click **New**.
- 5 Enter the general settings for the new BGP neighbor.
  - a Enter an IPv4 or IPv6 address for the new BGP neighbor.
  - b Enter a remote Autonomous System (AS) number in ASPLAIN format.
  - c Enter a time interval between sending keep-alive messages to a BGP peer.
  - d Enter a time interval before declaring a BGP peer dead.
  - e From the drop-down menu, select a **Graceful Restart Mode** option for this neighbor.

Option	Description
<b>Disable</b>	Overrides the global edge gateway settings and deactivates graceful restart mode for this neighbor.
<b>Helper only</b>	Overrides the global edge gateway settings and configures graceful restart mode as <b>Helper only</b> for this neighbor.
<b>Graceful restart and Helper</b>	Overrides the global edge gateway settings and configures graceful restart mode as <b>Graceful restart and Helper</b> for this neighbor.

- f Toggle on the **AllowAS-in** toggle to enable receiving routes with the same AS.
- g If the BGP neighbor requires authentication, enter the password for the BGP neighbor.

- 6 Configure the Bidirectional Forwarding Detection (BFD) settings for the new BGP neighbor.
  - a (Optional) Toggle on the **BFD** option to enable BFD for failure detection.
  - b In the BDF interval text box, define the time interval for sending heartbeat packets.
  - c In the **Dead Multiple** text box, enter the number of times the BGP neighbor can fail to send heartbeat packets before the BFD declares it is down.
- 7 (Optional) Configure route filtering.
  - a From the **IP Address Family** drop-down menu, select an IP address family.
  - b To configure an inbound filter, select an IP prefix list.
  - c To configure an outbound filter, select an IP prefix list.
- 8 Click **Save**.

#### What to do next

You can view the status of each BGP neighbor, edit, or delete BGP neighbors as needed.

## Working with NSX Advanced Load Balancing

As an **organization administrator**, by configuring virtual services which distribute traffic across multiple server pools, you can balance the workloads in your data centers that are backed by NSX-T Data Center.

Starting with version 10.2, VMware Cloud Director provides load-balancing services by using the capabilities of VMware NSX Advanced Load Balancer (Avi Networks).

VMware Cloud Director supports L4 and L7 load balancing that you can configure on an NSX-T Data Center edge gateway.

Level 4 load balancing (L4) directs traffic based on data from network and transport layer protocols, such as IP address and TCP port.

Level 7 load balancing (L7) distributes traffic based on attributes such as HTTP header, uniform resource identifier, SSL session ID, and HTML form data.

### Enable Load Balancer on an NSX-T Data Center Edge Gateway

Before an **organization administrator** can configure load balancing services, a **system administrator** must enable the load balancer on the NSX-T Data Center edge gateway.

#### Prerequisites

- Verify that you are a **system administrator**.
- Verify that you integrated VMware NSX Advanced Load Balancer in your cloud infrastructure. For more information on managing NSX Advanced Load Balancer, see *VMware Cloud Director Service Provider Admin Portal Guide*.

**Procedure**

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the NSX-T Data Center edge gateway on which you want to enable load balancing.
- 3 Under Load Balancer, click **General Settings**.
- 4 Click **Edit** and toggle on the **Load Balancer State** option.
- 5 Enter a network CIDR for a service network subnet from which to use IP addresses for creation of virtual services.

You can use the default service network subnet, by selecting the **Use Default** check box.

- 6 Click **Save**.

**What to do next**

[Assign a Service Engine Group to an NSX-T Data Center Edge Gateway.](#)

**Assign a Service Engine Group to an NSX-T Data Center Edge Gateway**

Before an **organization administrator** can configure load balancing services on an NSX-T Data Center edge gateway, a **system administrator** must assign a service engine group to the edge gateway.

The load balancing compute infrastructure provided by NSX Advanced Load Balancer is organized into service engine groups. A **system administrator** can assign one or more service engine groups to an NSX-T Data Center edge gateway.

All service engine groups that are assigned to a single edge gateway use the same service network.

**Prerequisites**

- Verify that you are a **system administrator**.
- [Enable Load Balancer on an NSX-T Data Center Edge Gateway.](#)

**Procedure**

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the NSX-T Data Center edge gateway to which you want to assign a service engine group.
- 3 Under Load Balancer, click **Service Engine Groups**.
- 4 Click **Add**.
- 5 Select an available service engine group from the list.
- 6 Enter a number for the maximum number of virtual services that can be placed on the edge gateway.
- 7 Enter a number for the guaranteed virtual services available to the edge gateway.

8 To confirm your settings, click **Save**.

## Edit the Settings of a Service Engine Group

A **system administrator** can edit the maximum number of supported virtual services and the number of reserved virtual services for a service engine group.

After you sync a service engine group, if the new maximum number of supported virtual services is lower than the number of reserved virtual services, the service engine group is marked as overallocated.

If a service engine group is overallocated, the creation of a new virtual service might fail, even if the edge gateway on which you create the virtual service has enough reserved capacity.

To avoid failure of virtual service creation, when you edit the settings of a service engine group, do not reduce the maximum number of supported virtual services below the number of initially reserved virtual services.

### Prerequisites

- Verify that you are a **system administrator**.
- [Enable Load Balancer on an NSX-T Data Center Edge Gateway](#).
- [Assign a Service Engine Group to an NSX-T Data Center Edge Gateway](#).

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the NSX-T Data Center edge gateway to which the service engine group is assigned.
- 3 Under Load Balancer, click **Service Engine Groups**.
- 4 Click **Edit**.
- 5 Edit the number for the maximum allowed virtual services that the edge gateway can use.  
Do not reduce the number unless mandatory. Otherwise, you might face failures when you create virtual services.
- 6 Edit the number for the guaranteed virtual services available to the edge gateway.
- 7 Click **Save**.

## Add a Load Balancer Server Pool

A server pool is a group of one or more servers that you configure to run the same application and to provide high availability.

### Prerequisites

- Verify that you are an **organization administrator**.
- Verify that your **system administrator** has enabled load balancing on the NSX-T edge gateway.

- Verify that your **system administrator** has assigned at least one service engine group to the edge gateway.

#### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the NSX-T Data Center edge gateway for which you want to configure a load balancer pool.
- 3 Under Load Balancer, click **Pools**, and then click **Add**.
- 4 Configure the general settings for the load balancer pool.
  - a Enter a meaningful name and, optionally, a description for the server pool.
  - b Select an algorithm balancing method.

The load balancing algorithm defines how incoming connections are distributed among the members of the server pool.

Option	Description
<b>Least Connections</b>	New connections are sent to the server that currently has the fewest connections.
<b>Round Robin</b>	New connections are sent to the next eligible server in the pool in a sequential order.
<b>Fastest Response</b>	New connections are sent to the server that provides the fastest response to new connections or requests.
<b>Consistent Hash</b>	New connections are distributed across the servers by using the IP address of the client to generate an IP hash key.
<b>Least Load</b>	New connections are sent to the server with the lightest load, regardless of the number of connections that server has.
<b>Fewest Servers</b>	Instead of attempting to distribute all connections or requests across all servers, the load balancer determines the fewest number of servers required to satisfy the current client load.
<b>Random</b>	The load balancer picks servers at random.
<b>Fewest Tasks</b>	Load is adaptively balanced, based on the server feedback.
<b>Core Affinity</b>	Each CPU core uses a subset of servers, and each server is used by a subset of cores. Essentially, it provides a many-to-many mapping between servers and cores.

- c To enable the server pool upon creation, toggle on the **State** option.
- d Enter a default destination server port to be used for the traffic to the pool member.
- e (Optional) In the **Graceful Disable Timeout** text box, enter the maximum time in minutes to deactivate gracefully a pool member.

The virtual service waits for the specified time before closing the existing connections to deactivated members.

- f (Optional) To activate a passive health monitor, toggle on the **Passive Health Monitor** option.
- g (Optional) Select an active health monitor.

Option	Description
HTTP	An HTTP request and response are used to validate the health.
HTTPS	Used against HTTPS encrypted web servers to validate the health.
TCP	A TCP connection is used to validate the health.
UDP	A UDP datagram is used to validate the health.
PING	An ICMP ping is used to validate the health.

5 Add a member to the server pool.

- a Click the **Members** tab and click **Add**.
- b Enter an IP address for the pool member.
- c Toggle on the **State** option to enable the pool member.
- d (Optional) Add a custom port for the server pool member.

The port number defaults to the destination port that you entered for the pool.

- e Enter a ratio for the pool member.

The ratio of each pool member denotes the traffic that goes to each server pool member. A server with a ratio of 2 gets twice as much traffic as a server with a ratio of 1. The default value is 1.

6 On the **SSL Settings** tab, configure the SSL settings for validating the certificates presented by the members of the load balancer pool.

- a To activate SSL, toggle on the **SSL Enable** option.
- b To hide certificates with private keys and see a list of CA certificates only, select the **Hide service certificates** check box.

7 To activate common name check for server certificates, toggle on the **Common Name Check** option and enter up to 10 domain names for the pool.

8 Click **Save**.

**What to do next**

[Create a Virtual Service.](#)

## Create a Virtual Service

A virtual service listens for traffic to an IP address, processes client requests, and directs valid requests to a member of the load balancer server pool.

A virtual service is a combination of an IP address and a port that uses a single network protocol. The virtual service is advertised to outside networks and is listening for client requests. When a client connects to the virtual service, the load balancer directs the request to a member of the load balancer server pool that you configured.

To secure SSL termination for a virtual service, you can use a certificate from the certificate library. For more information, see [Import Certificates to the Certificates Library](#).

### Prerequisites

- Verify that you are an **organization administrator**.
- Verify that your **system administrator** has enabled load balancing on the NSX-T edge gateway.
- Verify that your **system administrator** has assigned at least one service engine group to the edge gateway.
- [Add a Load Balancer Server Pool](#).

### Procedure

- 1 In the top navigation bar, click **Networking** and click the **Edge Gateways** tab.
- 2 Click the NSX-T Data Center edge gateway on which you want to create a virtual service.
- 3 Under Load Balancer, click **Virtual Services**, and then click **Add**.
- 4 Enter a meaningful name and, optionally, a description, for the virtual service.
- 5 To activate the virtual service upon creation, toggle on the **Enabled** option.
- 6 Select a service engine group for the virtual service.
- 7 Select a load balancer pool for the virtual service.
- 8 Enter an IP address for the virtual service.
- 9 Select the virtual service type.

Option	Description
HTTP	The virtual service listens for non-secure layer 7 HTTP requests. When you select this service type, it autopopulates the service port text box to 80, which you can replace with another valid port number.
HTTPS	The virtual service listens for secure level 7 HTTPS requests. When you select this service type, it autopopulates the service port text box to port 443, which you can replace with another valid port number. Select an SSL certificate to be used for SSL termination.

Option	Description
L4	The virtual service listens for layer 4 requests. When you select this service type, it autopopulates the service port text box to 80, which you can replace with another valid port number.
L4 TLS	The virtual service listens for secure layer 4 TLS requests. When you select this service type, it autopopulates the service port text box to TCP port 443, which you can replace with another valid port number. Select an SSL certificate to be used for SSL termination.

10 Click **Save**.

# Using Named Disks and Reviewing Storage Policies

# 6

You can create and manage named disks, and review the organization virtual data center storage policies by using the VMware Cloud Director tenant portal.

This chapter includes the following topics:

- [Creating and Using Named Disks](#)
- [Review Storage Policy Properties](#)

## Creating and Using Named Disks

Named disks are standalone virtual disks that you create in Organization VDCs. **Organization administrators** and users who have the respective rights can create, remove, and update named disks, and connect them to virtual machines.

When you create a named disk, it is associated with an Organization VDC but not with a virtual machine. After you create the disk in a VDC, the disk owner or an administrator can attach it to any virtual machine deployed in the VDC. If you have the **Create a Shared Disk** right, you can create a shared named disk that you can attach to multiple VMs. The disk owner can also modify the disk properties, detach it from a virtual machine, and remove it from the VDC. **System administrators** and **organization administrators** have the same rights to use and modify the disk as the disk owner.

---

**Note** Although vSphere supports configurations like Windows Server Failover Cluster (WSFC) and you can create a shared disk through physical SCSI bus sharing, VMware Cloud Director 10.2 does not support this feature. When creating a shared disk in VMware Cloud Director, you only create an underlying independent persistent disk in vSphere with multiwriter mode enabled.

---

If you attach a named disk, you cannot take VM snapshots. If a shared disk is attached to a VM, you cannot edit its hard disk setting from the VM details view.

If the organization VDC has a storage policy with enabled VM encryption, you can encrypt VMs and disks by associating them with storage policies that have the VM Encryption capability. See [Virtual Machine Encryption](#).

## Create a Named Disk

You can create a named disk and attach it to one or more virtual machines at a later stage.

To create a named disk, you must specify its name and size. You can optionally include a description and select a storage profile to be used by the disk. You can create a shared disk that you can attach to multiple VMs.

---

**Note** Although vSphere supports configurations like Windows Server Failover Cluster (WSFC) and you can create a shared disk through physical SCSI bus sharing, VMware Cloud Director 10.2 does not support this feature. When creating a shared disk in VMware Cloud Director, you only create an underlying independent persistent disk in vSphere with multiwriter mode enabled.

---

#### Prerequisites

- 1 You must have an **organization administrator** role or disk owner rights.
- 2 If you want to create a shared disk, you must have the **Create a Shared Disk** right.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Storage**, from the left panel, select **Named Disks**.
- 2 Click **New**.
- 3 Enter a name and, optionally, a description of the disk.
- 4 Select the storage policy from the **Storage Policy** drop-down menu.
- 5 Enter the size of the named disk.
- 6 Select the bus type and subtype, from the **Bus Type** and **Bus Sub-Type** drop-down menus, respectively.
- 7 If you want to attach the named disk to multiple VMs, select the **Shareable** check box.  
You cannot edit this setting later.
- 8 Click **Save**.

## Edit a Named Disk

After you have created the disk, you can modify its name, description, storage policy, and size.

You cannot edit the **Shareable** setting of a named disk.

#### Prerequisites

- 1 You must have an **organization administrator** role or disk owner rights.

#### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Storage**, from the left panel, select **Named Disks**.
- 2 Select the disk you want to modify, and click **Edit**.
- 3 Edit the settings such as name, description, storage policy, and size.

- 4 Click **Save**.

## Attach a Named Disk to a Virtual Machine

After you create a named disk in a VDC, you can attach it to any virtual machine that is deployed in the VDC. You can attach a shared named disk to multiple VMs.

### Prerequisites

You must have an **organization administrator** role or disk owner rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Storage**, from the left panel, select **Named Disks**.
- 2 Click the radio button next to the name of the named disk that you want to attach to a virtual machine, and click **Attach**.
- 3 From the drop-down menu, select a virtual machine to which to attach the named disk, and click **Apply**.
- 4 If you want to attach another VM to a shared disk, repeat [Step 2](#) and [Step 3](#).

### What to do next

You can attach more named disks to the VM or detach them as needed.

## Delete a Named Disk

If you don't need a named disk, you can delete it.

### Prerequisites

You must have an **organization administrator** role or disk owner rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore and under **Storage**, from the left panel, select **Named Disks**.
- 2 Select the disk you want to delete, and click **Delete**.
- 3 Click **OK**.

## Review Storage Policy Properties

You can review the storage policies and storage policy details.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

## Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore.
- 2 Under **Storage**, click **Storage Policies**.  
The list of the available storage policies displays.
- 3 To view the details about a storage policy, click the name of the storage policy.
- 4 Review the details on the **General** and **Metadata** tabs, and click **OK**.  
You can review the name, limit, IOPS settings, and metadata details of the storage policy.

# Reviewing and Editing Virtual Data Center Properties

# 7

As an **organization administrator**, you can review the virtual data center properties. You can also control the access to organization VDCs by users and groups in your organization.

This chapter includes the following topics:

- [Review Virtual Data Center Properties](#)
- [Review the Virtual Data Center Metadata](#)
- [Limit Access to an Organization VDC to Specific Users and Groups in Your Organization](#)

## Review Virtual Data Center Properties

You can review the properties of the virtual data centers that are assigned to your organization.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore.
- 2 Under **Settings**, click **General**.

### Results

You can review the properties of the virtual data center, such as name, description, and status. Metrics information about the data center includes the allocation model and vCPU, as well as CPU, and memory usage.

## Review the Virtual Data Center Metadata

VMware Cloud Director provides a general-purpose facility for associating user-defined metadata with an object. If your system administrator has created metadata for the organization virtual data center, you can review the organization data center metadata.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center you want to explore.
- 2 Under **Settings**, click **Metadata**.  
The list of the available metadata displays.

## Limit Access to an Organization VDC to Specific Users and Groups in Your Organization

As an **organization administrator**, you can limit the access to each of the organization VDCs in your organization to specific users and groups.

By default, organization VDCs are shared with all users and groups that have a role which includes the **Allow Access to All Organization VDCs** right.

If your organization has multiple organization VDCs and you want to have them managed separately, you can create a custom role that would function as an organization VDC administrator and assign it to specific users or groups within your organization, providing them with access only to a specific VDC's compute and networking resources.

### Prerequisites

- 1 Verify that you are an **organization administrator**.
- 2 Create a custom role for the users and groups that you want to provide with access to a specific organization VDC. This role must exclude the **Allow Access to All Organization VDCs** right. See [Chapter 13 Managing Users, Groups and Roles](#).

### Procedure

- 1 On the **Virtual Data Center** dashboard screen, click the card of the virtual data center that you want to limit access to.
- 2 Under **Settings**, click **Sharing**.  
The list of users and groups within the organization that have access to the VDC appears.
- 3 To change the access settings to the organization VDC, click **Edit**.
- 4 Select **Specific Users and Groups**.
- 5 From the **Users** list, select the users that you want to provide with access to the VDC.
- 6 From the **Groups** list, select the groups that you want to provide with access to the VDC.
- 7 To share the VDC with the selected users and groups, click **Share**.

## Results

Access to the organization VDC is limited to the users and groups that you selected.

# Working with Dedicated vCenter Server Instances, Endpoints, and Proxies



You can access a dedicated vCenter Server environment or vCenter Server components from the VMware Cloud Director Tenant Portal.

## Dedicated vSphere Data Centers

In VMware Cloud Director, a Software-Defined Data Center (SDDC) encapsulates an entire dedicated vCenter Server environment.

Dedicated vCenter Server instances in VMware Cloud Director remove the requirement for a vCenter Server instance to be publicly accessible.

The **system administrator** can publish one or more dedicated vCenter Server instances to your organization. You can use the endpoints to access the UI or API of proxied or non-proxied components.

## Endpoints

A dedicated vCenter Server instance can include one or more endpoints that provide access to different components from the underlying environment. Endpoints can provide an access point to a data center component, such as a vCenter Server instance, an ESXi host, an NSX Manager instance, or an NSX-T Manager instance.

Endpoints might or might not be connected to a proxy.

## Proxies

VMware Cloud Director can act as an HTTPS proxy server and provide access to a dedicated vCenter Server instance, and to different components of shared or dedicated vCenter Server instances that are backing up your environment.

You can log in to the UI or API of the proxied components by using your VMware Cloud Director account.

To access proxied components, you must either use Chrome Browser Extension for VMware Cloud Director, or manually configure your browser with your proxy Settings.

This chapter includes the following topics:

- [Using Chrome Browser Extension for VMware Cloud Director](#)
- [Configure Your Browser with Your Proxy Settings](#)
- [Log In to the UI of a Component by Using an Endpoint](#)

## Using Chrome Browser Extension for VMware Cloud Director

You can use Chrome Browser Extension for VMware Cloud Director to log in to the proxied vSphere components in your environment.

Chrome Browser Extension for VMware Cloud Director provides proxy configuration and authentication.

Chrome Browser Extension for VMware Cloud Director supports multisite environments.

You can add the extension to your Chrome browser through the [Chrome Web Store](#).

## Configure Your Browser with Your Proxy Settings

Before you can access the UI of a proxied vSphere component, you must set up the proxies that are published to your organization.

To configure your browser to use your published proxies, you copy the URL of the proxy auto-config (PAC) file into your browser.

---

**Note** When the **system administrator** publishes a dedicated vSphere data center to your organization, or adds a proxy to one of your dedicated vSphere data centers, it is possible that it takes a few minutes for the browser to refetch the PAC from the provided URL. To force a refresh of the browser, you can repeat this procedure.

---

### Prerequisites

- Verify that the **system administrator** published at least one dedicated and enabled vCenter Server instance to your organization.
- Verify that the **system administrator** published the **SDDC\_VIEW** and **Token: Manage** rights to your organization, and your role includes these rights.
- Verify that the **system administrator** published and enabled the **CPOM extension** plug-in to your organization. This plug-in provides the function for viewing and using dedicated vSphere data centers in the VMware Cloud Director Tenant Portal.

### Procedure

- 1 In the top navigation bar, click **Data Centers** and then click **Virtual Data Center**.
- 2 On the **Dedicated vSphere Data Centers** pane, click **Click here to view Proxy Configuration Guide**.
- 3 Copy the PAC URL and click **Next**.

- 4 Follow the instructions to configure your browser to point to the PAC URL.
- 5 If a proxied component is using self-signed certificates, import the certificates to your browser.
  - a On the target vSphere data center card, click **Actions**, and click **Import Certificate**.
  - b Download the certificate and the certificate revocation list (CRL).
  - c Import the downloaded certificate to your browser.See the user instructions for your browser.

## Log In to the UI of a Component by Using an Endpoint

You can use endpoints to access the UI of proxied or non-proxied components with your VMware Cloud Director account.

### Prerequisites

If you want to access a proxied component, [Configure Your Browser with Your Proxy Settings](#), or [Using Chrome Browser Extension for VMware Cloud Director to Google Chrome](#).

### Procedure

- 1 In the top navigation bar, click **Data Centers** and then click **Virtual Data Center**.
- 2 Select the **Dedicated vSphere Data Centers** tab.
- 3 Open the endpoint of the dedicated vCenter Server instance.
  - To open the default endpoint, click **Open vSphere**.
  - To open a non-default endpoint, follow these steps:
    - Click the **Actions** menu, and click **View Endpoints**.
    - Click the endpoint URL.

If you are accessing a proxied component, a new card with your proxy credentials opens.

- 4 If you are logging in a proxied component, access the component by using your credentials.
  - a Copy the user name and the password.
  - b To activate the proxy, click **Open**.

A new card opens and prompts you for authentication against the proxy.
  - c In the **User Name** text box, paste the copied user name.
  - d In the **Password** text box, paste the copied password and click **OK**.

# Working with vApp Templates

# 9

A vApp template is a virtual machine image that is loaded with an operating system, applications, and data. These templates ensure that virtual machines are consistently configured across an entire organization. vApp templates are added to catalogs.

This chapter includes the following topics:

- [View a vApp Template](#)
- [Create a vApp Template from an OVF File](#)
- [Import a Virtual Machine from vCenter Server as a vApp Template](#)
- [Assign a VM Placement Policy and a VM Sizing Policy to a vApp Template](#)
- [Download a vApp Template](#)
- [Delete a vApp Template](#)

## View a vApp Template

You can see the list of vApp templates that are available in the catalogs, to which you have access. You can view a vApp template and explore the virtual machines that it contains.

You can access only vApp templates that are included in catalogs items that have been shared to you. For more information about sharing catalogs, see [Share a Catalog](#).

### Prerequisites

This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.

The list of templates appears in a grid view.

- 2 (Optional) Configure the grid view to contain elements you want to see.
  - a From the grid view, click the grid editor icon (  ) below the list of vApp templates.
  - b Select the elements you want to include in the grid view, such as version, status, catalog, owner, and so on.
  - c Click **OK**.

The grid displays the elements you selected for each vApp template in the list.
- 3 To view the virtual machines included in a vApp template, click the vApp template name.
 

The virtual machines that the vApp template includes display in a grid.
- 4 (Optional) To select the elements you want to see in the grid view, click the grid editor icon (  ) below the list of virtual machines.
  - a Select the elements you want to include in the grid view.
  - b Click **OK**.

## Create a vApp Template from an OVF File

You can upload an OVF package to create a vApp template in a catalog.

VMware Cloud Director supports the Open Virtualization Format (OVF) and Open Virtualization Appliance (OVA) specifications. If you upload an OVF file that includes OVF properties for customizing its virtual machines, those properties are preserved in the vApp template. For information about creating OVF packages, see the *OVF Tool User Guide* and *VMware vCenter Converter User's Guide*.

### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.
 

The list of templates appears in a grid view.
- 2 Click **New**.
- 3 Enter a URL address of the OVF file, or click the **Upload** icon to browse to a location accessible from your computer and select the OVF/OVA template file.
 

The location might be your local hard drive, a network share, or a CD/DVD drive. The supported file extensions include `.ova`, `.ovf`, `.vmdk`, `.mf`, `.cert`, and `.strings`. If you select to upload an OVF file, which references more files than you are trying to upload, for example, a VMDK file, you must browse and select all files.
- 4 Verify the details of the OVF/OVA template you are about to deploy and click **Next**.

- 5 Enter a name and, optionally, a description for the vApp template, and click **Next**.
- 6 From the **Catalog** drop-down menu, select the catalog, to which you want to add the template.
- 7 Review the vApp template settings, and click **Finish**.

### Results

The new vApp template appears in the templates grid view.

## Import a Virtual Machine from vCenter Server as a vApp Template

If you have **system administrator** rights, you can import vCenter Server VMs to VMware Cloud Director as vApp templates in catalogs.

### Prerequisites

To see and import VMs from vCenter Server as vApp templates, verify that you have **system administrator** rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.  
The list of templates appears in a grid view.
- 2 Click **Import from vCenter**.
- 3 From the drop-down menu, select a vCenter Server instance from which to import the vApp template.
- 4 Select a template from the list of virtual machines.
- 5 Enter a name and, optionally, a description for the vApp template.
- 6 From the drop-down menu, select a catalog to which to add the vApp template.
- 7 (Optional) To delete the source virtual machine, toggle on the **Move Virtual Machine** option.
- 8 (Optional) Mark the vApp template as a preferred template in the catalog.
- 9 Click **Import**.

## Assign a VM Placement Policy and a VM Sizing Policy to a vApp Template

To associate the VMs of a vApp template with specific VM placement and VM sizing policies, you can tag individual VMs of a vApp template with the policies you want to assign.

Starting with VMware Cloud Director 10.0, you can allow the users to change the predefined VM placement or VM sizing policies while editing a VM.

---

**Note** After you upgrade to VMware Cloud Director 10.0 or later, all preexisting template taggings become modifiable. If you want to disallow the changes to the predefined VM placement or VM sizing policies, you must deselect the **Modifiable** check box for the policies that you want to be unchangeable.

---

#### Prerequisites

- This operation requires the right to edit a vApp template.
- Verify that you have at least one vApp template in your VMware Cloud Director environment.

#### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.  
The list of templates appears in a grid view.
- 2 Select the radio button next to the vApp template you want to tag, and click **Tag with Compute Policies**.
- 3 If you want to assign a VM placement policy to a VM in the vApp template, select a policy from the **VM Placement Policy** drop-down menu on the row corresponding to the VM.
- 4 If you want to assign a VM sizing policy to a VM in the vApp template, select a policy from the **VM Sizing Policy** drop-down menu on the row corresponding to the VM.
- 5 (Optional) To allow the users to change the predefined VM placement or VM sizing policies while editing a VM, select the **Modifiable** check box under the policy drop-down menu.
- 6 Click **Tag**.

## Download a vApp Template

You can download a vApp template from a catalog as an OVA file to your local machine.

#### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

#### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.  
The list of templates appears in a grid view.

- 2 Click the list bar (  ) on the left of the vApp template you want to download, and select **Download**.

---

**Note** You can download vApp templates from your organization catalogs. If you are an organization administrator, you can download vApp templates from a public catalog. Otherwise, the **Download** button is dimmed.

---

- 3 (Optional) To preserve the UUIDs and MAC addresses of the virtual machines in the downloaded OVA package, select the **Preserve identity information** check box.
- 4 Click **OK** and wait for the download to complete.

The OVA file is saved to the default download location of your Web browser.

## Delete a vApp Template

You can delete a vApp template from an organization catalog. If the catalog is published, the vApp template is also deleted from public catalogs.

### Prerequisites

This operation requires the rights included in the predefined **vApp Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **vApp Templates**.  
The list of templates appears in a grid view.
- 2 Click the list bar (  ) on the left of the vApp template you want to delete, and select **Delete**.
- 3 Confirm the deletion.

The deleted vApp template is removed from the grid view.

# Working with Media Files

# 10

The catalog allows you to upload, copy, move, and edit the properties of media files.

This chapter includes the following topics:

- [Upload Media Files](#)
- [Delete a Media File](#)
- [Download a Media File](#)

## Upload Media Files

You can upload new media files or new versions of existing media files to a catalog. Users with access to the catalog can open the media files with their virtual machines.

### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Media & Other**.  
The list of media files appears in a grid view.
- 2 Click **Add**.
- 3 From the **Catalog** drop-down menu, select a catalog to which you want to upload the media file.
- 4 Enter a name for the media file.  
If you do not enter a name, the name text box is populated automatically after the name of the media file.
- 5 Click the upload icon to browse and select the disk image file, for example an `.iso` file.
- 6 Click **OK**.  
After the upload starts, the media file appears in the grid.

### What to do next

Depending on the file size, it might take some time for the upload to complete. You can monitor the status of the upload in the **Recent Tasks** view. For more information, see [View Tasks](#).

## Delete a Media File

You can delete media files that you no longer want to use from your catalog.

### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Media & Other**.  
The list of media files appears in a grid view.
- 2 Click the list bar (  ) on the left of the media file you want to delete, and select **Delete**.
- 3 Confirm the deletion.  
The deleted media file is removed from the grid view.

## Download a Media File

You can download a media file from a catalog.

### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Media & Other**.  
The list of media files appears in a grid view.
- 2 Click the list bar (  ) on the left of the media file you want to download, and select **Download**.  
The download task starts, and the file is saved to the default download location of your web browser.

### What to do next

Depending on the file size, it might take some time for the download to complete. You can monitor the status of the download in the **Recent Tasks** panel. For more information, see [View Tasks](#).

# Working with Catalogs

# 11

A catalog is a container for vApp templates and media files in an organization. Organization administrators and catalog authors can create catalogs in an organization. Catalog contents can be shared with other users or organizations in the VMware Cloud Director installation or published externally for access by organizations outside the VMware Cloud Director installation.

VMware Cloud Director contains private catalogs, shared catalogs, and externally accessible catalogs. Private catalogs include vApp templates and media files that you can share with other users in the organization. If a system administrator enables catalog sharing for your organization, you can share an organization catalog to create a catalog accessible to other organizations in the VMware Cloud Director installation. If a system administrator enables external catalog publishing for your organization, you can publish an organization catalog for access by organizations outside the VMware Cloud Director installation. An organization outside the VMware Cloud Director installation must subscribe to an externally published catalog to access its contents.

You can upload an OVF package directly to a catalog, save a vApp as a vApp template, or import a vApp template from vSphere. See [Create a vApp Template from an OVF File](#) and [Save a vApp as a vApp Template to a Catalog](#).

Members of an organization can access vApp templates and media files that they own or that are shared with them. Organization administrators and system administrators can share a catalog with everyone in an organization or with specific users and groups in an organization. See [Share a Catalog](#).

This chapter includes the following topics:

- [View Catalogs](#)
- [Create a Catalog](#)
- [Share a Catalog](#)
- [Delete a Catalog](#)
- [Change the Owner of a Catalog](#)
- [Manage Metadata for a Catalog](#)
- [Publish a Catalog](#)
- [Subscribe to an External Catalog](#)
- [Update the Location URL and the Password for a Subscribed Catalog](#)

- [Synchronize a Subscribed Catalog](#)

## View Catalogs

You can access catalogs shared with you within your organization. You can access public catalogs if an organization administrator has made them accessible within your organization.

Catalog access is controlled by catalog sharing, not by the rights in your role. You can access only those catalogs or catalog items that are shared with you. For more information, see [Share a Catalog](#).

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

- 2 (Optional) Configure the grid view to contain elements you want to see.

- a From the grid view, click the grid editor icon (  ) displayed below the list of catalogs.
- b Select the elements you want to include in the grid view, such as version, description, status, and so on.
- c Click **OK**.

The grid displays the elements you selected for each catalog.

- 3 (Optional) From the grid view, use the list bar (  ) to display the actions you can take for each catalog.

For example, you can share or delete a catalog.

## Create a Catalog

You can create new catalogs and associate them with a storage policy.

### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

- 2 Click **New** to create a new catalog.
- 3 Enter the name and, optionally, a description of the catalog.

- 4 (Optional) Select whether you want to assign a storage policy to the catalog, and select a storage policy.
- 5 Click **OK**.

### Results

The new catalog appears in the grid view on the **Catalogs** tab.

## Share a Catalog

You can share a catalog with all members of your organization, or with specific members.

### Prerequisites

- This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.
- You must be the owner of the catalog.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

- 2 Click the list bar (  ) on the left of the catalog you want to share, and select **Share**.

The list of users who can access the catalog appears in the grid view of the **Share Catalog** window.

- 3 Click **Add** to share the catalog with other users.

Option	Description
Share with everyone in this organization	Grant access to all users and groups in the organization.
Share with specific users and groups	Select the users or groups to whom you want to grant catalog access, and click <b>Add</b> .

- 4 Select the access level.

Option	Description
Read Only	Users with access to this catalog have read access to the vApp templates and ISO files of the catalog.
Read/Write	Users with access to this catalog have read access to the vApp templates and ISO files of the catalog and can add vApp templates and ISO files to the catalog.
Full Control	Users with access to this catalog have full control of the contents and settings of the catalog.

5 Click **OK**.

The users or groups that now have access to the catalog appear in the grid view of the **Share Catalog** dialog box.

6 (Optional) Select to share read-only access to the administrators of all other organizations

7 Click **Save**.

### Results

On the **Catalogs** tab, the Shared status for this catalog in the grid view changes.

## Delete a Catalog

You can delete a catalog from your organization.

### Prerequisites

This operation requires the rights included in the predefined **Catalog Author** role or an equivalent set of rights.

---

**Note** The catalog must not contain any vApp templates or media files. You can move these items to a different catalog or delete them.

---

### Procedure

1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

2 Click the list bar (  ) on the left of the catalog you want to delete, and select **Delete**.

3 Confirm the deletion.

The deleted catalog item is removed from the grid view.

## Change the Owner of a Catalog

An **organization administrator** can change the owner of a catalog.

Before you can delete a user who owns a catalog, you must change the owner or delete the catalog.

### Prerequisites

This operation requires the rights included in the predefined **organization administrator** role or an equivalent set of rights.

### Procedure

1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

- 2 Click the list bar (  ) on the left of a catalog, and select **Change owner**.

The list of users who can access the catalog appears in the grid view of the **Change Owner** window.

- 3 Select the user that you want to make the new owner of the catalog and click **OK**.

#### Results

On the **Catalogs** tab, the name of the owner of the catalog in the grid view changes.

## Manage Metadata for a Catalog

As an **organization administrator** or a **catalog owner**, you can create or update the metadata for the catalogs that you own.

#### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

- 2 Click the list bar (  ) on the left of a catalog, and select **Metadata**.

The metadata for the selected catalog is displayed in a grid view.

- 3 (Optional) To add metadata, click **Add**.

- a Enter the metadata name.

The name must be unique within the metadata names attached to this object.

- b Select the metadata type, such as **Text**, **Number**, **Date and Time**, or **Yes or No**.

- c Enter the metadata value.

- d Click **Save**.

- 4 (Optional) Update existing metadata.

You cannot update the metadata name.

- a Update the metadata type.

- b Enter the new metadata value.

- c Click **Save**.

- 5 (Optional) Delete existing metadata.

- a Click the delete icon.

- b Click **Save**.

## Publish a Catalog

If the **system administrator** has granted you catalog access, you can publish a catalog externally to make its vApp templates and media files available for subscription by organizations outside the VMware Cloud Director installation.

### Prerequisites

Verify that the **system administrator** enabled external catalog publishing for the organization and granted you catalog access.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.  
The list of catalogs appears in a grid view.
- 2 Click the list bar (  ) on the left of the catalog you want to publish, and select **Publish Settings**.
- 3 Select **Enable Publishing** and, optionally, enter a password for catalog access.  
Only ASCII characters are supported.
- 4 Click **Save**.

## Subscribe to an External Catalog

You can subscribe to an external catalog and thus create a read-only copy of an externally published catalog. You cannot modify a subscribed catalog.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- The **system administrator** must grant your organization permission to subscribe to external catalogs.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.  
The list of catalogs appears in a grid view.
- 2 Click **New** to create a new catalog.
- 3 Enter the name and, optionally, a description of the catalog.
- 4 Select to subscribe to an external catalog and provide the subscription URL.
- 5 Enter the optional password to access the catalog.
- 6 Select whether you want to automatically download the content from the external catalog.

- 7 Click **OK**.

## Update the Location URL and the Password for a Subscribed Catalog

After you create a subscribed catalog, you can update the location URL and the password for the subscribed catalog.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- You must have created a subscribed catalog.
- The **system administrator** must grant your organization permission to subscribe to external catalogs.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.  
The list of catalogs appears in a grid view.
- 2 Click the list bar (  ) on the left of a subscribed catalog, and select **Subscribe settings**.  
If the catalog is not a subscribed one, the option is dimmed.
- 3 Update the location URL and the password for this subscribed catalog.
- 4 Select whether you want to download the content from the external catalog automatically.
- 5 Click **Save**.

## Synchronize a Subscribed Catalog

After you create a subscribed catalog, you can synchronize it with the original catalog to see if there are any changes. For example, if the metadata of the original catalog is changes, when you perform the synchronization, the metadata of the subscribed catalog is updated.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- You must have created a subscribed catalog.
- The **system administrator** must grant your organization permission to subscribe to external catalogs.

## Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Catalogs**.

The list of catalogs appears in a grid view.

- 2 Click the list bar (  ) on the left of a subscribed catalog, and select **Sync**.

If the catalog is not a subscribed one, the option is dimmed.

The subscribed catalog is synchronized with the original one.

# Working with Organization Virtual Data Center Templates

# 12

As an organization administrator or any role that has rights to view and instantiate organization virtual data center templates, you can create additional organization virtual data centers.

An organization virtual data center template specifies a configuration for an organization virtual data center and, optionally, an Edge Gateway, and organization virtual data center network. System administrators can enable organization administrators to create these resources in their organizations by creating organization virtual data center templates and sharing them with those organizations.

By creating and sharing virtual data center templates, system administrators enable self-service provisioning of organization virtual data centers while retaining administrative control over allocation of system resources, such as provider virtual data centers and external networks.

System administrators create organization virtual data center templates and provide different organizations with access to the templates.

If your organization has been provided with access to virtual data center templates, you can use the VMware Cloud Director Tenant Portal to create virtual data centers from the available templates.

This chapter includes the following topics:

- [View Available Virtual Data Center Templates](#)
- [Instantiate a Virtual Data Center from a Template](#)

## View Available Virtual Data Center Templates

You can view the organization virtual data center templates that a system administrator has created for you.

View the virtual data center templates before you create a new organization virtual data center from the virtual data center template.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or a role that has rights to view and instantiate organization virtual data center templates.

### Procedure

- ◆ In the top navigation bar, click **Libraries** and in the left panel, select **Organization VDC Templates**.

The list of virtual data center templates appears in a grid view.

### What to do next

Review the descriptions of the organization virtual data center templates and select the template from which you want to create a new organization virtual data center.

## Instantiate a Virtual Data Center from a Template

When a system administrator creates an organization virtual data center (VDC) template and publishes the template to your organization, you can create an organization VDC from the template.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or a role that has rights to view and instantiate organization VDC templates.

### Procedure

- 1 In the top navigation bar, click **Libraries** and in the left panel, select **Organization VDC Templates**.

The list of virtual data center templates appears in a grid view.

- 2 Select a template, and click **New VDC**.

Starting with VMware Cloud Director 10.2.2, after selecting a template, you must click **Instantiate VDC**.

- 3 Enter a name of the VDC and, optionally, a description.
- 4 Click **Create**.

### Results

The creation of the new organization virtual data center is instantiated and might take a few minutes. You can see the progress of the task in the **Recent Tasks** panel.

### What to do next

You can manage your newly created organization virtual data center by creating virtual machines, vApps, managing the network and security settings, and so on.

# Managing Users, Groups and Roles

# 13

You can add organization administrators to VMware Cloud Director individually, or as part of an LDAP group. You can also add and modify the roles that determine what rights a user has within their organization.

---

**Important** You must be an **organization administrator** to manage the users, groups, and roles within your organization. Your **system administrator** can publish one or more global tenant roles to your tenant, and as an **organization administrator**, you can see them in the list of roles. Such roles are for example, **Catalog Author**, **vApp Author**, **vApp User**, **Organization Administrator**, and so on. You cannot modify the predefined global tenant roles, but you can create and update similar custom tenant roles and assign them to users within your tenant.

---

This chapter includes the following topics:

- [Managing Users](#)
- [Managing Groups](#)
- [Roles and Rights](#)

## Managing Users

From the tenant portal you can create, edit, import, and delete users. In addition, you can also unlock user accounts in case a user tried to log in with an incorrect password and as a result has locked their own user account.

### Create a User

You can create a user within your VMware Cloud Director organization.

#### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

#### Procedure

- 1 In the top navigation bar, click **Administration**.

- 2 In the left panel, under **Access Control**, click **Users**.

The list of users appears.

- 3 Click **New**.

- 4 Enter a user name and the password setting of the user.

The minimum password length is six characters.

- 5 Select whether to enable the user upon creation.

- 6 If you want to set a specific limitation on the resources available to the user, turn on the **Configure user's quota** toggle.

If you turn on the toggle, when you complete this wizard, VMware Cloud Director redirects you to the **Quotas** page. You can add quotas on the number of Tanzu Kubernetes clusters, all or running VMs managed by the user, consumed CPU, memory, and storage. Select **Unlimited** if you want the user to have unlimited resources of the selected type.

- 7 Choose the role that you want to assign to the user.

The **Available roles** menu consist of a list of predefined roles and any custom roles that you or the system administrator might have created.

Predefined role	Description
<b>vApp Author</b>	The rights associated with the predefined <b>vApp Author</b> role allow a user to use catalogs and create vApps.
<b>Console Access Only</b>	The rights associated with the predefined <b>Console Access Only</b> role allow a user to view virtual machine state and properties and to use the guest OS.
<b>vApp User</b>	The rights associated with the predefined <b>vApp User</b> role allow a user to use existing vApps.
<b>Organization Administrator</b>	A user with the predefined <b>Organization Administrator</b> role can use the VMware Cloud Director tenant portal or the Cloud Director OpenAPI to manage users and groups in their organization and assign them roles, including the predefined <b>Organization Administrator</b> role. An <b>organization administrator</b> can use the Cloud Director OpenAPI to create or update role objects that are local to the organization. Roles created or modified by an <b>organization administrator</b> are not visible to other organizations.
<b>Defer to Identity Provider</b>	Rights associated with the predefined <b>Defer to Identity Provider</b> role are determined based on information received from the user's OAuth or SAML Identity Provider. To qualify for inclusion when a user is assigned the <b>Defer to Identity Provider</b> role, a role name supplied by the Identity Provider must be an exact, case-sensitive match for a role, or name defined in your organization.
<b>Catalog Author</b>	The rights associated with the predefined <b>Catalog Author</b> role allow a user to create and publish catalogs.

- 8 (Optional) Enter the contact information, such as name, email address, phone number, and instant messaging ID.

- 9 Click **Save**.

## What to do next

If you enabled quotas configuration for the user and VMware Cloud Director redirects you to the **Quotas** page, see [Manage the Resource Quotas of a User](#).

## Import Users

You can add users to your organizations by importing an LDAP user or a SAML user and assigning them a certain role.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- Verify that you have a valid connection to an LDAP server or that you [Enable Your Organization to Use a SAML Identity Provider](#).

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel, under **Access Control**, click **Users**.

The list of users appears.

- 3 Click **Import Users**.
- 4 Select a source from which you want to import the users.

You will only view the source LDAP server or SAML server that you configured as identity provider.

Source	Action
LDAP	Import users from an LDAP server. <ol style="list-style-type: none"> <li>a Enter a full or partial name in the text box and click <b>Search</b>.</li> <li>b Select the users whom you want to import and click <b>Add</b>.</li> </ol>
SAML	Import users from a SAML server. Enter the user names of the users that you want to import. User names must be in the name identifier format supported by the SAML identity provider configured for this organization. <p><b>Note</b> If you are using vCenter Single Sign-On as the SAML identity provider, the user names that you import from a vCenter Single Sign-On domain must be in User Principal Name (UPN) format, for example jdoe@mydomain.com.</p> Use a new line for each user name.

- 5 Select the role which you want to assign to the users that you import.
- 6 Click **Save**.

## Modify a User

As an organization administrator, you can modify the password, the contact, and the virtual machine quota settings of an existing user. In addition, you can also change the role of the user.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel, under **Access Control**, click **Users**.  
The list of users appears.
- 3 Click the radio button next to the name of the user that you want to edit and click **Modify**.
- 4 Update the settings you want to modify.
  - a Change the password as necessary.
  - b Select whether to activate or deactivate the user.
  - c Update the user role.
  - d Update the contact information, such as name, email address, phone number, and instant messaging ID.
  - e Edit virtual machine quota for the user.
- 5 Click **Save**.

## Deactivate or Activate a User Account

You can deactivate a user account to prevent that user from logging in to VMware Cloud Director. To delete a user, you must first deactivate their account.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel, under **Access Control**, click **Users**.  
The list of users appears.
- 3 To deactivate a user account, click the radio button next to the user name, click **Disable**, and confirm.

- 4 To activate a user account that you have already deactivated, click the radio button next to the user name, and click **Enable**.

## Delete a User

You can remove a user from the VMware Cloud Director organization by deleting the user account.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- Deactivate the account you want to delete.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel, under **Access Control**, click **Users**.  
The list of users appears.
- 3 Click the radio button next to the name of the user that you want to delete and click **Delete**.
- 4 To confirm that you want to delete the user account, click **OK**.

## Unlock a Locked Out User Account

In case you have enabled a lockout policy in your VMware Cloud Director organization, a user account is locked after a certain number of invalid login attempts. You can unlock the locked user account. Best practice is to change the password of the user and unlock the account.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel, under **Access Control**, click **Users**.  
The list of users appears.
- 3 Click the radio button next to the user name, click **Unlock**.

## Manage the Resource Quotas of a User

You can manage the overall resource consumption limit of a user. You can add, edit, and remove the user's quotas on VMs, Tanzu Kubernetes clusters, CPU, memory, or storage.

Users can see the quotas relevant only to their user type. Users inherit quotas from the group they belong to. If a user inherits a resource quota from their group and has an explicit user-level quota defined for that resource, then the user-level quota takes priority over the group-level quota.

For information about creating or importing users, see [Create a User](#) or [Import Users](#).

### Prerequisites

Verify that you have the necessary rights to add, edit, and delete resource quotas. By default, **Organization administrators** can change the quotas of users.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Users**.
- 3 Select the name of a user and select the **Quotas** tab.

Users do not have any quotas by default. All users that belong to a group inherit the group's quotas. If the user belongs to a group that has a quota on resources, the quota appears in the user's list of quotas as not editable.

- 4 Click **Edit**.
- 5 Modify the quota for the selected user.

You can add, edit, or remove quotas on the number of Tanzu Kubernetes clusters, all or running VMs managed by the user, consumed CPU, memory, and storage. Select **Unlimited** if you want the user to have unlimited resources of the selected type.

- 6 Click **Save**.

## Managing Groups

If you have a valid connection to an LDAP server or have enabled your organization to use a SAML identity provider, you can import an LDAP group or a SAML group. You can also edit or delete an imported group.

### Import a Group

To add a group of users, you can import an LDAP group or a SAML group.

#### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- Verify that you have a valid connection to an LDAP server or that you [Enable Your Organization to Use a SAML Identity Provider](#).

**Procedure**

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Groups**.

The list of user groups appears.

- 3 Click **Import Group**.
- 4 Select a source from which you want to import the user group.

You can only view the source LDAP server or SAML server that you configured as identity provider.

Source	Action
LDAP	Import a user group from an LDAP server. <ol style="list-style-type: none"> <li>a Enter a full or partial name in the text box and click <b>Search</b>.</li> <li>b Select the user groups that you want to import and click <b>Add</b>.</li> </ol>
SAML	Import user groups from a SAML server. Enter the names of the groups that you want to import. Use a new line for each group name.

- 5 Select the role which you want to assign to the group of users that you import.
- 6 Click **Save**.

**What to do next**

If you enabled quotas configuration for the group and VMware Cloud Director redirects you to the **Quotas** page, see [Manage the Resource Quotas of a Group](#).

**Delete a Group**

You can remove a group from your VMware Cloud Director organization by deleting their LDAP group.

When you delete an LDAP group, users who have a VMware Cloud Director account based solely on their membership in that group are stranded and cannot log in.

**Prerequisites**

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

**Procedure**

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Groups**.

The list of user groups appears.

- 3 Click the radio button next to the name of the group that you want to delete, and click **Delete**.
- 4 To confirm that you want to delete the group, click **OK**.

## Edit a Group

You can edit a group from the VMware Cloud Director tenant portal.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Groups**.  
The list of user groups appears.
- 3 Click the radio button next to the name of the group that you want to delete, and click **Edit**.
- 4 Edit the group as necessary.
  - a Change the description.
  - b Change the role of the members of the group as necessary.
- 5 Click **Save**.

## Manage the Resource Quotas of a Group

By directly setting quota on a group, you can manage the overall resource consumption limit of each user in it. You can add, edit, and remove the group's quotas on VMs, Tanzu Kubernetes clusters, CPU, memory, or storage. Quotas of the group are applied on each member of the group.

Users inherit quotas from the group they belong to. If a user inherits a resource quota from their group and has an explicit user-level quota defined for that resource, then the user-level quota takes priority over the group-level quota.

For information about importing groups, see [Import a Group](#).

### Prerequisites

Verify that you have the necessary rights to add, edit, and delete resource quotas. By default, **Organization administrators** can change the quotas of groups.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Groups**.

- 3 Select the name of a group and select the **Quotas** tab.

Groups do not have any quotas by default. All users that belong to a group inherit the group's quotas. If the user belongs to a group that has a quota on resources, the quota appears in the user's list of quotas as not editable.

- 4 Click **Edit**.

- 5 Modify the quota for the selected group.

You can add, edit, or remove quotas on the number of Tanzu Kubernetes clusters, all or running VMs managed by the group, consumed CPU, memory, and storage. Select **Unlimited** if you want the group of users to have unlimited resources of the selected type.

- 6 Click **Save**.

## Roles and Rights

VMware Cloud Director uses roles and rights to determine what actions a user can perform in an organization. VMware Cloud Director includes a number of predefined roles with specific rights.

**System administrators** and **organization administrators** must assign each user or group a role. The same user can have a different role in different organizations. **System administrators** can create roles and modify existing ones for the whole system, while **organization administrators** can create and modify roles only for the organization that they administer.

The VMware Cloud Director tenant portal allows **organization administrators** to manage the roles in their organization. If a **system administrator** publishes one or more predefined tenant roles to your organization, as an **organization administrator** you can see these roles, but you cannot modify them. You can, however, create custom tenant roles with similar rights and assign them to the users within your organization.

For information about the predefined roles and their rights, see [Predefined Roles and Their Rights](#).

## Predefined Roles and Their Rights

Each VMware Cloud Director predefined role contains a default set of rights required to perform operations included in common workflows. By default, all predefined global tenant roles are published to every organization in the system.

### Predefined Provider Roles

By default, the provider roles that are local only to the provider organization are the **System Administrator** and **Multisite System** roles. **System administrators** can create additional custom provider roles.

#### System Administrator

The **System Administrator** role exists only in the provider organization. The **System Administrator** role includes all rights in the system. For a list of rights available only to

the **System administrator** role, see the *VMware Cloud Director Service Provider Admin Portal Guide*. The **System administrator** credentials are established during installation and configuration. A **System Administrator** can create additional system administrator and user accounts in the provider organization.

### Multisite System

Used for running the heartbeat process for multisite deployments. This role has only a single right, **Multisite: System Operations**, which gives a permission to make a Cloud Director OpenAPI request that retrieves the status of the remote member of a site association.

### Predefined Global Tenant Roles

By default, the predefined global tenant roles and the rights they contain are published to all organizations. **System Administrators** can unpublish rights and global tenant roles from individual organizations. **System Administrators** can edit or delete predefined global tenant roles. **System administrators** can create and publish additional global tenant roles.

### Organization Administrator

After creating an organization, a **System Administrator** can assign the role of **Organization Administrator** to any user in the organization. A user with the predefined **Organization Administrator** role can manage users and groups in their organization and assign them roles, including the predefined **Organization Administrator** role. Roles created or modified by an **Organization Administrator** are not visible to other organizations.

### Catalog Author

The rights associated with the predefined **Catalog Author** role allow a user to create and publish catalogs.

### vApp Author

The rights associated with the predefined **vApp Author** role allow a user to use catalogs and create vApps.

### vApp User

The rights associated with the predefined **vApp User** role allow a user to use existing vApps.

### Console Access Only

The rights associated with the predefined **Console Access Only** role allow a user to view virtual machine state and properties and to use the guest OS.

### Defer to Identity Provider

Rights associated with the predefined **Defer to Identity Provider** role are determined based on information received from the user's OAuth or SAML Identity Provider. To qualify for inclusion when a user or group is assigned the **Defer to Identity Provider** role, a role or group name supplied by the Identity Provider must be an exact, case-sensitive match for a role or group name defined in your organization.

- If an OAuth Identity Provider defines the user, the user is assigned the roles named in the `roles` array of the user's OAuth token.
- If a SAML Identity Provider defines the user, the user is assigned the roles named in the SAML attribute whose name appears in the `RoleAttributeName` element, which is in the `SamlAttributeMapping` element in the organization's `OrgFederationSettings`.

If a user is assigned the **Defer to Identity Provider** role but no matching role or group name is available in your organization, the user can log in to the organization but has no rights. If an Identity Provider associates a user with a system-level role such as **System Administrator**, the user can log in to the organization but has no rights. You must manually assign a role to such users.

Except the **Defer to Identity Provider** role, each predefined role includes a set of default rights. Only a **System Administrator** can modify the rights in a predefined role. If a **System administrator** modifies a predefined role, the modifications propagate to all instances of the role in the system.

## Rights in Predefined Global Tenant Roles

Various rights are common to multiple predefined global roles. These rights are granted by default to all new organizations, and are available for use in other roles created by the **Organization Administrator**. For a list of the rights in predefined tenant roles, see [Rights in Predefined Global Tenant Roles](#).

## Rights in Predefined Global Tenant Roles

Various rights are common to multiple predefined global roles. These rights are granted by default to all new organizations, and are available for use in other roles created by the **Organization Administrator**.

## Rights Included in the Global Tenant Roles in VMware Cloud Director

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Access All Organization VDCs	✓				
	Catalog: Add vApp from My Cloud	✓	✓	✓		
	Catalog: Change Owner	✓				
	Catalog: CLSP Publish Subscribe	✓	✓			

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Catalog: Create / Delete a Catalog	✓	✓			
	Catalog: Edit Properties	✓	✓			
	Catalog: Publish	✓	✓			
	Catalog: Sharing	✓	✓			
	Catalog: View ACL	✓	✓			
	Catalog: View Private and Shared Catalogs	✓	✓	✓		
	Catalog: View Published Catalogs	✓				
	Custom entity: View all custom entity instances in org	✓				
	Custom entity: View custom entity instance	✓				
	Disk: Change Owner	✓	✓			
	Disk: Create	✓	✓	✓		
	Disk: Delete	✓	✓	✓		
	Disk: Edit Properties	✓	✓	✓		
	Disk: View Encryption Status	✓		✓		
	Disk: View Properties	✓	✓	✓	✓	
	General: Administrator Control	✓				
	General: Administrator View	✓				
	General: Send Notification	✓				
	Group / User: View	✓				
	Hybrid Cloud Operations: Acquire control ticket	✓				
	Hybrid Cloud Operations: Acquire from-the-cloud tunnel ticket	✓				
	Hybrid Cloud Operations: Acquire to-the-cloud tunnel ticket	✓				
	Hybrid Cloud Operations: Create from-the-cloud tunnel	✓				

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Hybrid Cloud Operations: Create to-the-cloud tunnel	✓				
	Hybrid Cloud Operations: Delete from-the-cloud tunnel	✓				
	Hybrid Cloud Operations: Delete to-the-cloud tunnel	✓				
	Hybrid Cloud Operations: Update from-the-cloud tunnel endpoint tag	✓				
	Hybrid Cloud Operations: View from-the-cloud tunnel	✓				
	Hybrid Cloud Operations: View to-the-cloud tunnel	✓				
	Organization Network: Edit Properties	✓				
	Organization Network: View	✓				
	Organization vDC Compute Policy: View	✓	✓	✓	✓	
	Organization vDC Distributed Firewall: Configure Rules	✓				
	Organization vDC Distributed Firewall: View Rules	✓				
	Organization vDC Gateway: Configure DHCP	✓				
	Organization vDC Gateway: Configure DNS	✓				
	Organization vDC Gateway: Configure ECMP Routing	✓				
	Organization vDC Gateway: Configure Firewall	✓				
	Organization vDC Gateway: Configure IPSec VPN	✓				
	Organization vDC Gateway: Configure Load Balancer	✓				
	Organization vDC Gateway: Configure NAT	✓				
	Organization vDC Gateway: Configure Static Routing	✓				
	Organization vDC Gateway: Configure Syslog	✓				

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Organization vDC Gateway: Convert to Advanced Networking	✓				
	Organization vDC Gateway: View	✓				
	Organization vDC Gateway: View DHCP	✓				
	Organization vDC Gateway: View DNS	✓				
	Organization vDC Gateway: View Firewall	✓				
	Organization vDC Gateway: View IPSec VPN	✓				
	Organization vDC Gateway: View Load Balancer	✓				
	Organization vDC Gateway: View NAT	✓				
	Organization vDC Gateway: View Static Routing	✓				
	Organization vDC Network: Edit Properties	✓				
	Organization vDC Network: View Properties	✓		✓		
	Organization vDC Storage Policy: View Capabilities	✓				
	Organization vDC Storage Profile: Set Default	✓				
	Organization vDC: Edit	✓				
	Organization vDC: Edit ACL	✓				
	Organization vDC: Manage Firewall	✓				
	Organization vDC: View	✓	✓			
	Organization vDC: View ACL	✓				
	Organization VDC: view metrics	✓				
	Organization vDC: VM-VM Affinity Edit	✓	✓	✓		
	Organization: Edit Association Settings	✓				
	Organization: Edit Federation Settings	✓				

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Organization: Edit LDAP Settings	✓				
	Organization: Edit Leases Policy	✓				
	Organization: Edit OAuth Settings	✓				
	Organization: Edit Password Policy	✓				
	Organization: Edit Properties	✓				
	Organization: Edit Quotas Policy	✓				
	Organization: Edit SMTP Settings	✓				
	Organization: Import User/Group from IdP while Editing VDC ACL	✓				
	Organization: View	✓	✓	✓		
	Organization: view metrics	✓				
✓	Quota Policy Capabilities: View	✓				
	Role: Create, Edit, Delete, or Copy	✓				
	Service Library: View service libraries	✓				
	UI Plugins: View	✓	✓	✓	✓	
	vApp Template / Media: Copy	✓	✓	✓		
	vApp Template / Media: Create / Upload	✓	✓			
	vApp Template / Media: Edit	✓	✓	✓		
	vApp Template / Media: View	✓	✓	✓	✓	
	vApp Template: Change Owner	✓	✓			
	vApp Template: Checkout	✓	✓	✓	✓	
	vApp Template: Download	✓	✓			
	vApp: Change Owner	✓				
	vApp: Copy	✓	✓	✓	✓	
	vApp: Create / Reconfigure	✓	✓	✓		
	vApp: Delete	✓	✓	✓	✓	
	vApp: Download	✓	✓	✓		

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	vApp: Edit Properties	✓	✓	✓	✓	
	vApp: Edit VM Compute Policy	✓	✓	✓		
	vApp: Edit VM CPU	✓	✓	✓		
	vApp: Edit VM Hard Disk	✓	✓	✓		
	vApp: Edit VM Memory	✓	✓	✓		
	vApp: Edit VM Network	✓	✓	✓	✓	
	vApp: Edit VM Properties	✓	✓	✓	✓	
	vApp: Manage VM Password Settings	✓	✓	✓	✓	✓
	vApp: Power Operations	✓	✓	✓	✓	
	vApp: Sharing	✓	✓	✓	✓	
	vApp: Snapshot Operations	✓	✓	✓	✓	
	vApp: Upload	✓	✓	✓		
	vApp: Use Console	✓	✓	✓	✓	✓
	vApp: View ACL	✓	✓	✓	✓	
	vApp: View VM and VM's Disks Encryption Status	✓		✓		
	vApp: View VM metrics	✓		✓	✓	
	vApp: VM Boot Options	✓	✓	✓		
	vApp: VM Metadata to vCenter	✓	✓	✓		
✓	VDC Group: Configure	✓				
✓	VDC Group: View	✓				
✓	VDC Group: Configure Logging	✓				
	VDC Template: Instantiate	✓				
	VDC Template: View	✓				

## Create a Custom Tenant Role

Organization administrators can use the tenant portal to create custom tenant role objects in the organizations they administer.

## Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

## Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Roles**.

The list of roles appears.

- 3 Click **Add**.
- 4 Enter a name and, optionally, a description of the role.
- 5 Expand the rights for the role and select the rights for the role.

The rights are grouped in categories and subcategories that allow either viewing or managing objects.

Option	Description
<b>Access Control</b>	Rights controlling the access to view and manage certain objects.
<b>Administration</b>	Rights controlling the administrative access.
<b>Compute</b>	Rights controlling access and management of the organization and provider virtual data centers, the vApps, organization virtual data centers templates, virtual machine groups, and virtual machine monitoring.
<b>Extensions</b>	Rights controlling the access to any additional plug-ins and VMware Cloud Director extensions.
<b>Infrastructure</b>	Rights controlling the access and management of the infrastructure objects, such as datastores, disks, hosts, and so on.
<b>Libraries</b>	Rights controlling access and management of any catalogs and catalog items.
<b>Networking</b>	Rights controlling access and management of the network settings.

- 6 Click **Save**.

## Edit a Custom Tenant Role

Organization administrators can use the tenant portal to edit custom tenant role objects in the organizations they administer. As an organization administrator you can only view the global tenant roles that a system administrator has published to your organization. You cannot edit global tenant roles.

## Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Roles**.  
The list of roles appears.
- 3 Click the radio button next to the role that you want to edit, and click **Edit**.
- 4 Modify the role settings as needed.
  - a Change the name and, optionally, the description of the role.
  - b Edit the rights for the role.
- 5 Click **Save**.

## Delete a Role

Organization administrators can use the tenant portal to delete role objects in the organizations they administer.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel under **Access Control**, click **Roles**.  
The list of roles appears.
- 3 Click the radio button next to the role that you want to delete, and click **Delete**.
- 4 Confirm that you want to delete the role by clicking **OK**.

# Configure Identity Providers

# 14

You can integrate your cloud with an external identity provider and import users and groups to your organization.

You can enable your organization to use a SAML identity provider or you can configure an LDAP server connection.

This chapter includes the following topics:

- [Enable Your Organization to Use a SAML Identity Provider](#)
- [Edit LDAP Settings for Your Organization](#)
- [Configure, Test, and Synchronize an LDAP Connection](#)

## Enable Your Organization to Use a SAML Identity Provider

Enable your organization to use a Security Assertion Markup Language (SAML) identity provider, also called single sign-on, to import users and groups from a SAML identity provider and allow imported users to sign on to the organization with the credentials established in the SAML identity provider.

When you import users and groups, the system extracts a list of attributes from the SAML token, if available, and uses them for interpreting the corresponding pieces of information about the user attempting to log in.

- `email address = "EmailAddress"`
- `user name = "UserName"`
- `full name = "FullName"`
- `user's groups = "Groups"`
- `user's roles = "Roles"`

The role attribute is configurable.

Group information is necessary if the user is not directly imported but is expected to be able to log in by virtue of membership in imported groups. A user might belong to multiple groups, and can have multiple roles during a session.

If an imported user or group is assigned the **Defer to Identity Provider** role, the roles are assigned based on the information gathered from the Roles attribute in the token. If a different attribute is used, this attribute name can be configured by using the API only, and only the Roles attribute is configurable. If the **Defer to Identity Provider** role is used, but no role information can be extracted, the user can log in but does not have any rights to perform any activities.

### Prerequisites

- This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.
- Verify that you have access to an SAML 2.0 compliant identity provider.
- Verify that you receive the required metadata from your SAML identity provider. You must import the metadata to VMware Cloud Director either manually or as an XML file. The metadata must include the following information:
  - The location of the single sign-on service
  - The location of the single logout service
  - The location of the service's X.509 certificate

For information on configuring and acquiring metadata from a SAML provider, see the documentation for your SAML identity provider.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 Under **Identity Providers**, click **SAML**.
- 3 Click **Edit**.
- 4 On the **Service Provider** tab, enter the Entity ID.

The Entity ID is the unique identifier of your organization to your identity provider. You can use the name of your organization, or any other string that satisfies the requirements of your SAML identity provider.

---

**Important** Once you specify an Entity ID, you cannot delete it. To change the Entity ID, you must do a full SAML reconfiguration for your organization. For information about Entity IDs, see [Assertions and Protocols for the OASIS Security Assertion Markup Language \(SAML\) 2.0](#).

---

- 5 Click the **Metadata** link to download the SAML metadata for your organization.  
The downloaded metadata must be provided as-is to your identity provider.
- 6 Review the Certificate Expiration date and, optionally, click Regenerate to regenerate the certificate used to sign federation messages.

The certificate is included in the SAML metadata, and is used for both encryption and signing. Either or both encryption and signing might be required depending on how trust is established between your organization and your SAML identity provider.

- 7 On the **Identity Provider** tab, enable the **Use SAML Identity Provider** toggle.
- 8 Copy and paste the SAML metadata you received from your identity provider to the text box, or click **Upload** to browse to and upload the metadata from an XML file.
- 9 Click **Save**.

#### What to do next

- Configure your SAML provider with VMware Cloud Director metadata. See your SAML identity provider documentation and the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.
- Import users and groups from your SAML identity provider. See [Chapter 13 Managing Users, Groups and Roles](#)

## Edit LDAP Settings for Your Organization

You can configure your organization to use the system LDAP connection as a shared source of users and groups. You can configure your organization to use a separate LDAP connection as a private source of users and groups.

#### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

#### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 In the left panel, under **Identity Providers**, click **LDAP**.

The current LDAP settings are displayed.

- 3 On the **LDAP Settings** tab, click **Edit**.
- 4 Configure the LDAP source of users and groups for your organization and click **Save**.

Option	Description
Do not use LDAP	The organization does not use an LDAP server as a source of organization users and groups.
VMware Cloud Director system LDAP service	The organization uses the VMware Cloud Director system LDAP connection configured by your service provider. Enter the distinguished name for the organizational unit.
Custom LDAP service	The organization uses a private LDAP server as a source of organization users and groups.

## What to do next

If you selected **Custom LDAP service**, click the **Custom LDAP** tab to [Configure, Test, and Synchronize an LDAP Connection](#).

# Configure, Test, and Synchronize an LDAP Connection

To configure an LDAP connection, you set the details of your LDAP server. You can test the connection to make sure that you entered the correct settings and the user and group attributes are mapped correctly. When you have a successful LDAP connection, you can synchronize the user and group information with the LDAP server at any time.

## Prerequisites

If you plan to connect to an LDAP server over SSL (LDAPS), verify that the certificate of your LDAP server is compliant with the Endpoint Identification introduced in Java 8 Update 181. The common name (CN) or the subject alternative name (SAN) of the certificate must match the FQDN of the LDAP server. For more information, see the *Java 8 Release Changes* at <https://www.java.com>.

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

## Procedure

- 1 In the **Connection** tab, enter the required information for the LDAP connection.

Required Information	Description
Server	The host name or IP address of the LDAP server.
Port	The port number on which the LDAP server is listening. For LDAP, the default port number is 389. For LDAPS, the default port number is 636.
Base distinguished name	The base distinguished name (DN) is the location in the LDAP directory where VMware Cloud Director to connect. To connect at root level, enter only the domain components, for example, <b>DC=example,DC=com</b> . To connect to a node in the domain tree structure, enter the distinguished name for that node, for example, <b>OU=ServiceDirector,DC=example,DC=com</b> . Connecting to a node limits the scope of the directory available to VMware Cloud Director.
Connector type	The type of your LDAP server. Can be <b>Active Directory</b> or <b>OpenLDAP</b> .
Use SSL	If your server is LDAPS, select this check box.
Accept all certificates	If your server is LDAPS, either select this check box or upload the LDAP SSL certificate.
Custom Truststore	If your server is LDAPS, either click the <b>Upload</b> button and import an LDAP SSL certificate or select <b>Accept all certificates</b> .

Required Information	Description
<b>Authentication method</b>	Simple authentication consists of sending the user's DN and password to the LDAP server. If you are using LDAP, the LDAP password is sent over the network in plain text.  If you want to use Kerberos, you must configure the LDAP connection by using the vCloud API.
<b>User name</b>	Enter the full LDAP distinguished name (DN) of a service account with domain admin rights. VMware Cloud Director uses this account to query the LDAP directory and retrieve user information.  If the anonymous read support is enabled on your LDAP server, you can leave these text boxes blank.
<b>Password</b>	The password for the service account that connects to the LDAP server.  If the anonymous read support is enabled on your LDAP server, you can leave these text boxes blank.

- 2 Click the **User Attributes** tab, examine the default values for the user attributes, and, if your LDAP directory uses different schema, modify the values.
- 3 Click the **Group Attributes** tab, examine the default values for the group attributes, and, if your LDAP directory uses different schema, modify the values.
- 4 Click **Save**.
- 5 If you selected the **Use SSL** check box, and if the certificate of the LDAPS server is not yet trusted, on the **Trust Certificate** window, confirm if you trust the certificate presented by the server endpoint.
- 6 To test the LDAP connection settings and the LDAP attribute mappings:
  - a Click **Test**
  - b Enter the password of the LDAP server user that you configured and click **Test**.  
  
If connected successfully, a green check mark is displayed.  
  
The retrieved user and group attribute values are displayed in a table. The values that are successfully mapped to LDAP attributes are marked with green check marks. The values that are not mapped LDAP attributes are blank and marked with red exclamation marks.
  - c To exit, click **Cancel**.
- 7 To synchronize VMware Cloud Director with the configured LDAP server, click **Sync**.  
  
VMware Cloud Director synchronizes the user and group information with the LDAP server regularly depending on the synchronization interval that you set in the general system settings.  
  
Wait a few minutes for the synchronization to finish.

## Results

You can import users and groups from the newly configured LDAP server.

You can import, download, edit, and delete certificates from VMware Cloud Director. You can copy the certificate PEM data to the clipboard.

This chapter includes the following topics:

- [Import Trusted Certificates](#)
- [Import Certificates to the Certificates Library](#)

## Import Trusted Certificates

You can import certificates of servers that VMware Cloud Director communicates with, such as vCenter Server, NSX Manager, and so on.

When using VMware Cloud Director in FIPS mode, you must use FIPS-compatible private keys. You can use pyOpenSSL to generate private keys in FIPS-compatible PKCS#8 format. If you generate PKCS#8 private keys by using OpenSSL, the private keys are not FIPS-compatible. For more information about FIPS mode, see [Activate FIPS Mode on the Cells in the Server Group](#) or [Activate or Deactivate FIPS Mode on the VMware Cloud Director Appliance](#).

### Prerequisites

Verify that you are logged in as a **system administrator** or **organization administrator**.

### Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Certificate Management**, select **Trusted Certificates** and click **Import**.
- 3 Upload a PEM file containing the certificates that you want to import and click **Import**.
- 4 (Optional) Edit the certificate name.
- 5 Click **Import**.

### What to do next

- Download a certificate.
- Edit a certificate name.
- Delete a certificate.

- Copy the PEM data to the clipboard.

## Import Certificates to the Certificates Library

In the VMware Cloud Director certificates library, you can import certificates used when creating entities that you must secure, such as servers, edge gateways, and so on.

The certificate library contains information about single certificates, certificate chains, private keys, certificate expiration dates, the entities that the certificates secure, and so on.

When using VMware Cloud Director in FIPS mode, you must use FIPS-compatible self-signed certificates and private keys. You can generate self-signed unencrypted certificates and private keys by using pyOpenSSL. If you generate self-signed certificates and private keys by using OpenSSL, the certificates and private keys are not FIPS-compatible. For more information about FIPS mode, see [Activate FIPS Mode on the Cells in the Server Group](#) or [Activate or Deactivate FIPS Mode on the VMware Cloud Director Appliance](#).

### Prerequisites

Verify that you are logged in as a **system administrator** or **organization administrator**.

### Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Certificate Management**, select **Certificates Library** and click **Import**.
- 3 Enter a name, and optionally, a description for this certificate in the certificate library and click **Next**.
- 4 Upload a PEM file containing the certificate chain that you want to import and click **Next**.
- 5 (Optional) Upload a private key file.  
Your private key file might not be protected with a passphrase.
- 6 Click **Import**.

### Results

The imported certificate appears in the list of available certificates during the creation of entities that you must secure.

### What to do next

- Download a certificate.
- Edit the name and description of a certificate.
- Delete a certificate. You can delete only certificates that do not secure any entities.
- Copy the certificate PEM data to the clipboard.

# Managing Your Organization

# 16

As an **organization administrator**, you can modify various settings within your organization. You can modify the name of the organization, email settings, domain settings, metadata, policies, and so on.

You can use the VMware Cloud Director API to subscribe to messages about events and tasks in your organization through the MQTT protocol. See the information about subscribing to events and tasks by using an MQTT client in the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

This chapter includes the following topics:

- [Edit the Organization Name and Description](#)
- [Modify Your Email Settings](#)
- [Test SMTP Settings](#)
- [Modify Domain Settings for the Virtual Machines in Your Organization](#)
- [Working with Multiple Sites](#)
- [Configure and Manage Multisite Deployments](#)
- [Understanding Leases](#)
- [Modify the vApp and vApp Template Lease Policies Within Your Organization](#)
- [Modify the Password and User Account Policies Within Your Organization](#)
- [Create an Advisories Dashboard](#)

## Edit the Organization Name and Description

You can edit the full name and the description of your organization.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.

- 2 Under **Settings**, click **General**.

The list of general settings, such as organization name, default URL, full name, and description displays.

- 3 To modify the full name and the description of the organization, click **Edit**.
- 4 Apply the necessary changes and click **Save**.

## Modify Your Email Settings

You can review and modify the default email settings that were set when the system administrator created your organization

VMware Cloud Director sends alert emails when having important information to report, for example, when a data store is running out of space. By default, an organization sends email alerts to the system administrators or a list of email addresses specified at the system level by using an SMTP server specified at the system level. You can modify the email settings at the organization level if you want VMware Cloud Director to send alerts for that organization to a different set of email addresses than those specified at the system level or you want the organization to use a different SMTP server to send alerts than the server specified at the system level.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 Under **Settings**, click **Email**.

The email settings for your organization are displayed.
- 3 Click **Edit**.
- 4 Edit the SMTP server settings on the **SMTP Server** tab.
  - a Select whether to use a custom SMTP server or the default.
  - b If you select to use a custom SMTP server, enter the DNS host name or IP address of the SMTP server in the **SMTP server name** text box.
  - c (Optional) Enter the SMTP server port.
  - d (Optional) Select whether to require authentication and enter a user name and password.
- 5 To edit the notification settings, click the **Notification Settings** tab.
  - a Select to use custom notification settings.
  - b Enter the email address that appears as the sender for organization emails.
  - c (Optional) Enter the text to use as the email subject prefix.

- d (Optional) Select whether to send notifications to all organization administrators or to specific email addresses.
  - e (Optional) If you select to send notifications to specific email addresses, enter the email addresses by separating them with a comma.
- 6 Click **Save**.

## Test SMTP Settings

After you modify the email settings for your organization, you can test the SMTP settings.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 Under **Settings**, click **Email**.  
The email settings for your organization are displayed.
- 3 Click **Test**.
- 4 Enter a destination email address and the SMTP server password to test the SMTP settings, and click the **Test** button.

## Modify Domain Settings for the Virtual Machines in Your Organization

You can set a default Windows domain which virtual machines created in your organization can join. Virtual machines can always join a domain for which they have credentials, regardless of whether you specify a default domain or not.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 Under **Settings**, click **Guest Personalization**.
- 3 Select to enable the domain join for the virtual machines in the organization.
- 4 Enter the domain name, user name, and password.  
The credentials that you enter apply to a regular domain user, not a domain administrator.

- 5 (Optional) Enter an account organizational unit.
- 6 Click **Save**.

## Working with Multiple Sites

The VMware Cloud Director Multisite feature enables a service provider or a tenant of multiple, geographically-distributed VMware Cloud Director installations (server groups) to manage and monitor those installations and their organizations as single entities.

The VMware Cloud Director Tenant Portal provides **organization administrators** with a way to associate organizations at associated sites.

For more information about site associations, see the *VMware Cloud Director Service Provider Admin Portal Guide*.

## Configure and Manage Multisite Deployments

After a **system administrator** has associated two sites, **organization administrators** at any member site can begin associating their organizations.

To create an association between two organizations (we will call them Org-A and Org-B here), you must be an **organization administrator** for both organizations so that you can log in to each organization, retrieve its local association data, and submit the retrieved data to the other organization.

---

**Important** The process of associating two organizations can be logically decomposed into two complementary pairing operations. The first operation (in this example) pairs Org-A at Site-A with Org-B at Site-B. You must then pair Org-B at Site-B with Org-A at Site-A. Until both pairings are complete, the association is incomplete.

---

### Prerequisites

- The sites occupied by the organizations must be associated.
- You must be a **system administrator** at both sites or an **organization administrator** at both organizations.

### Procedure

- 1 Log in to the VMware Cloud Director tenant portal of Org-A at Site-A to retrieve its local association data.
  - a Click **Administration**.
  - b Under **Settings**, click **Multisite**.
  - c To download the data in XML format, click **Export local association data**.

The browser saves the data in a file in its Downloads folder.

- 2 Log in to the VMware Cloud Director tenant portal of Org-B at Site-B to submit the local association data from Org-A at Site-A.

- a Click **Administration**.
- b Under **Settings**, click **Multisite**.
- c Click **Create new organization association**.

Submit the association data you downloaded in [Step 1](#) to Org-B by clicking the upload arrow below the **New Association XML** text box and selecting the local association data you downloaded in [Step 1](#).

- d Click **Next** to verify and submit the data.

The system pairs Org-A at Site-A with Org-B at Site-B.

- e Click **Finish** to view the associated organization.
- f To view details of the associated organization or delete the association, click the **Organization Name** card.

- 3 Complete the association by repeating Step 1 and Step 2 to retrieve the local association data from Org-B and submit it to Org-A.

## Understanding Leases

Creating an organization involves specifying leases. Leases provide a level of control over an organization's storage and compute resources by specifying the maximum amount of time that vApps can be running and that vApps and vApp templates can be stored.

The goal of a runtime lease is to prevent inactive vApps from consuming compute resources. For example, if a user starts a vApp and goes on a vacation without stopping it, the vApp continues to consume resources.

A runtime lease begins when a user starts a vApp. When a runtime lease expires, VMware Cloud Director stops the vApp.

The goal of a storage lease is to prevent unused vApps and vApp templates from consuming storage resources. A vApp storage lease begins when a user stops the vApp. Storage leases do not affect running vApps. A vApp template storage lease begins when a user adds the vApp template to a vApp, adds the vApp template to a workspace, downloads, copies, or moves the vApp template.

When a storage lease expires, VMware Cloud Director marks the vApp or vApp template as expired, or deletes the vApp or vApp template, depending on the organization policy you set.

# Modify the vApp and vApp Template Lease Policies Within Your Organization

You can review and modify the default policies that were set by the system administrator when your organization was created.

## Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

## Procedure

1 In the top navigation bar, click **Administration**.

2 Under **Settings**, click **Policies**.

You can view the default policies that your **system administrator** has set.

3 Click **Edit**.

4 Edit the vApp leases.

vApp leases provide a level of control over the organization storage and compute resources by specifying the maximum amount of time that vApps can be running and that vApps can be stored. You can also specify what happens to the vApps when their storage lease expires.

- a To define how long vApps can run before they are automatically stopped, enter the maximum runtime lease.
- b Select a runtime expiry action, such as power off or suspend.
- c To define how long stopped vApps remain available before being automatically cleaned up, enter the maximum storage lease.
- d Select a storage cleanup action, such as to delete permanently the vApps or move them to the expired items.

5 Edit the vApp template lease.

vApp template leases provide a level of control over the organization storage and compute resources by specifying the maximum amount of time that vApp templates can be stored. You can also specify what happens to the vApp templates when their storage lease expires.

- a To define how long the vApp templates remain available before being automatically cleaned up, enter the maximum storage lease.
- b Select a storage cleanup action, such as to delete permanently the vApp templates or move them to the expired items.

6 Click **OK**.

## Modify the Password and User Account Policies Within Your Organization

You can review and modify the default password and user account policies that were set by the system administrator when your organization was created.

The password and user account policies define the VMware Cloud Director behavior when a user enters an invalid password.

### Prerequisites

This operation requires the rights included in the predefined **Organization Administrator** role or an equivalent set of rights.

### Procedure

- 1 In the top navigation bar, click **Administration**.
- 2 Under **Settings**, click **Policies**.  
You can view the default policies that your **system administrator** has set.
- 3 Click **Edit**.
- 4 Enable locking of a user account after a number of invalid login attempts.
- 5 Enter the number of invalid login attempts before the account is locked.
- 6 Enter the time interval in minutes, in which the user with locked account cannot log back in.
- 7 Click **OK**.

## Create an Advisories Dashboard

You can create notifications that appear on top of the UI pages in the Tenant Portal. The messages can appear to the users within an organization or the users in all organizations.

You cannot edit advisories once you create them.

### Prerequisites

Verify that you are logged in as a **system administrator**.

### Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Settings**, select **Advisories** and click **New**.
- 3 In the description box, add the text of the notification.  
You can use basic Markdown to add links to the notifications.

#### 4 Select the priority of the message.

Different priority messages appear as different colors. The notifications appear in the order of their priority. Mandatory advisories cannot be dismissed or snoozed.

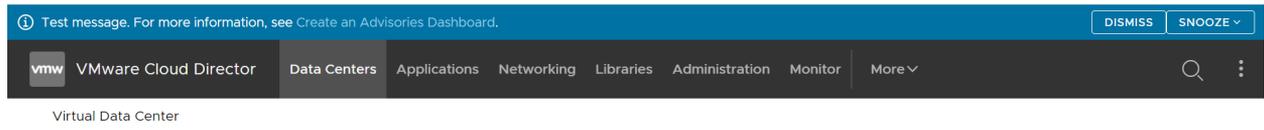
#### 5 Select the period for which you want the notification to appear in the UI.

You can view all advisories in the **Advisories** tab, however they appear to the selected group of users only during the selected period.

#### 6 Click **OK**.

### Results

The notification appears above the top navigation bar of the selected portal.



### What to do next

Delete the notification by selecting the radio button next to it and clicking **Delete**. The advisories appear in the **Advisories** tab even after they expire. To remove them from the list, you must delete them.

# Working with Service Library

# 17

The Service Library items in VMware Cloud Director are vRealize Orchestrator workflows that extend the cloud management capabilities and make it possible for administrators of either providers or tenants to monitor and manipulate different services.

This chapter includes the following topics:

- [Search for a Service](#)
- [Execute a Service](#)

## Search for a Service

The **Service Library** page in the VMware Cloud Director Tenant Portal lists the set of vRealize Orchestrator workflows that are imported to VMware Cloud Director and published to your organization.

### Prerequisites

This operation requires the Service Library rights to be included in the predefined user role.

### Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Service Library**.

The list of service items appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the service and a tag that corresponds to the service category where vRealize Orchestrator is imported.

- 2 In the **Search** text box on the top of the page, enter the first word of either the name of the service or the name of the category, to which the service belongs.
  - a Select whether you want to search among the names of the service or among the categories.

The search results display in a card view of twelve items per page, sorted by names in alphabetical order.

## Execute a Service

You can execute a service from the Service Library page in the VMware Cloud Director Tenant Portal.

### Prerequisites

This operation requires the Service Library rights to be included in the predefined user role.

### Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Service Library**.

The list of service items appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the service and a tag that corresponds to the service category where vRealize Orchestrator is imported.

- 2 Search for the service you want to execute.

- 3 Click **Execute** on the card of the service.

A new dialog opens. You must enter values for the required input parameters of the service.

- 4 Click **Finish** to confirm the execution of the service.

### What to do next

You can monitor the status of the execution in the **Recent Tasks** view. For more information, see [View Tasks](#).

Starting with VMware Cloud Director 10.2, service providers can use the VMware Cloud Director API to create extensions that provide additional VMware Cloud Director capabilities to the tenants. If a service provider granted you access, you can manage defined entities and share them with other tenants.

Service providers can create Runtime Defined Entities (RDEs) enabling extensions to store and manipulate the extension-specific information in VMware Cloud Director. For example, a Kubernetes extension can store information about the Kubernetes clusters it manages in RDEs. The extension can then provide extension APIs for managing those clusters using the information from the RDEs.

## Access to Defined Entities

Two complementary mechanisms control the access to RDEs.

- **Rights** - When a service provider creates an RDE type, they create a rights bundle for the type. A service provider must assign you one or more of the five type-specific rights: **View: TYPE**, **Edit: TYPE**, **Full Control: TYPE**, **Administrator View: TYPE**, and **Administrator Full Control: TYPE**.

The **View: TYPE**, **Edit: TYPE**, and **Full Control: TYPE** rights work only in combination with an ACL entry.

- **Access Control List (ACL)** - The ACL table contains entries defining the access users have to specific entities in the system. It provides an extra level of control over the entities. For example, while an **Edit: TYPE** right specifies that a user can modify entities to which they have access, the ACL table defines which entities the user has access to.

Table 18-1. Rights and ACL Entries for RDE Operations

Entity Operation	Option	Description
Read	<b>Administrator View: TYPE</b> right	Users with this right can see all RDEs of this type within an organization.
	<b>View: TYPE</b> right and ACL entry $\geq$ <b>View</b>	Users with this right and a read-level ACL can view RDEs of this type.

Table 18-1. Rights and ACL Entries for RDE Operations (continued)

Entity Operation	Option	Description
Modify	<b>Administrator Full Control: TYPE</b> right	Users with this right can create, view, modify, and delete RDEs of this type in all organizations.
	<b>Edit: TYPE</b> right and ACL entry $\geq$ <b>Change</b>	Users with this right and modify-level ACL can create, view, and modify RDEs of this type.
Delete	<b>Administrator Full Control: TYPE</b> right	Users with this right can create, view, modify, and delete RDEs of this type in all organizations.
	<b>Full Control: TYPE</b> right and ACL entry = <b>Full Control</b>	Users with this right and full control-level ACL can create, view, modify, and delete RDEs of this type.

## Sharing Defined Entities with Another User

If a **system administrator** published the rights bundle for a defined entity type and granted you `ReadWrite` or `FullControl` access or you are the defined entity owner, you can share the access to those entities with other users.

- 1 Assign the **View: TYPE**, **Edit: TYPE**, or **Full Control: TYPE** right from the bundle to the user roles you want to have the specific level of access to the defined entity.

---

**Note** You must be logged in as a **system administrator** or **organization administrator** to assign rights.

---

For example, if you want the users with the `tkg_viewer` role to view Tanzu Kubernetes clusters within the organization, you must add the **View: Tanzu Kubernetes Guest Cluster** right to the role. If you want the users with the `tkg_author` role to create, view, and modify Tanzu Kubernetes clusters within this organization, add the **Edit: Tanzu Kubernetes Guest Cluster** to that role. If you want the users with the `tkg_admin` role to create, view, modify, and delete Tanzu Kubernetes clusters within this organization, add the **Full Control: Tanzu Kubernetes Guest Cluster** right to the role.

- 2 Grant the specific user an Access Control List (ACL) by making the following REST API call.

```
POST https://[address]/cloudapi/1.0.0/entities/urn:vcloud:entity:[vendor]:[type name]:
[version]:[UUID]/accessControls
{
```

```

"grantType" : "MembershipAccessControlGrant",
"accessLevelId" : "urn:vcloud:accessLevel:[Access_level]",
"memberId" : "urn:vcloud:user:[User_ID]"
}

```

*Access\_level* must be `ReadOnly`, `ReadWrite`, or `FullControl`. *User\_ID* must be the ID of the user to which you want to grant the access to the defined entity.

You must have `ReadWrite` or `FullControl` access to an entity to grant ACL access to that entity.

Users with the **tkg\_viewer** role, described in the example, cannot grant ACL access. Users with the **tkg\_author** or **tkg\_admin** role can share access to a VMWARE:TKGCLUSTER entity with users who have the **tkg\_viewer**, **tkg\_author**, or **tkg\_admin** role by granting them ACL access using the API request.

Users with the **Administrator Full Control: Tanzu Kubernetes Guest Cluster** right can grant ACL access to any VMWARE:TKGCLUSTER entity.

You can also use REST API calls to revoke the access or to view who has access to the entity. See the VMware Cloud Director REST API documentation on [code.vmware.com](https://code.vmware.com).

## Changing the Owner of a Defined Entity

The owner of a defined entity or a user with the **Administrator Full Control: TYPE** right can transfer the ownership to another user by updating the defined entity model and changing the owner field with the ID of the new owner.

This chapter includes the following topics:

- [Working with Custom Entity Definitions](#)

## Working with Custom Entity Definitions

The custom entity definitions in VMware Cloud Director are object types that are bound to vRealize Orchestrator object types. Users within a VMware Cloud Director organization can own, manage, and change these types according to their needs. By executing services, organization users can instantiate the custom entities and apply actions over the instances of the objects.

### Search for a Custom Entity

You can search for those of the custom entities that were published to your organization.

#### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

## Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 In the **Search** text box on the top of the page, enter a word or a character of the name of the entity you want to find.

The search results display in a card view of twelve items per page, sorted by names in alphabetical order.

## Edit a Custom Entity Definition

You can modify the name and the description of a custom entity. You cannot change the type of the entity or the vRealize Orchestrator object type, to which the entity is bound, these are the default properties of the custom entity. If you want to modify any of the default properties, you must delete the custom entity definition and recreate it.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

## Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 In the card of the selected custom entity, select **Actions > Edit**.

A new dialog opens.

- 3 Modify the name or the description of the custom entity definition.

- 4 Click **OK** to confirm the change.

## Add a Custom Entity Definition

You can create a custom entity and map it to an existing vRealize Orchestrator object type.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

## Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 To add a new custom entity, click **New**.

A new dialog opens.

- 3 Follow the steps of the **Custom Entity Definition** wizard.

Step	
Name and Description	Enter a name and, optionally a description for the new entity. Enter a name for the entity type, for example <code>sshHost</code> .
vRO	From the drop-down menu, select the vRealize Orchestrator that you will use to map the custom entity definition.  <b>Note</b> If you have more than one vRealize Orchestrator server, you must create a custom entity definition for each one of them separately.
Type	Click the view list icon to browse through the available vRealize Orchestrator object types grouped by plug-ins. For example, <b>SSH &gt; Host</b> . If you know the name of the type, you can enter it directly in the text box. For example <code>SSH:Host</code> .
Review	Review the details that you specified and click <b>Done</b> to complete the creation.

## Results

The new custom entity definition appears in the card view.

## Custom Entity Instances

Running a vRealize Orchestrator workflow with an input parameter being an object type that is already defined as a custom entity definition in VMware Cloud Director shows the output parameter as an instance of a custom entity.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

## Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 In the card of the selected custom entity, click **Instances**.

The available instances display in a grid view.

- Click the list bar (  ) on the left of each entity to display the associated workflows.  
Clicking on a workflow initiates a workflow run which takes the entity instance as an input parameter.

## Associate an Action to a Custom Entity

By associating an action to a custom entity definition, you can execute a set of vRealize Orchestrator workflows on the instances of a particular custom entity.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

### Procedure

- In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.  
The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.
- In the card of the selected custom entity, select **Actions > Associate Action**.  
A new dialog opens.
- Follow the steps of the **Associate Custom Entity to VRO Workflow** wizard.

Step	Details
Select VRO Workflow	Select one of the listed workflows. These are the workflows that are available in the <b>Service Library</b> page.
Select Workflow Input Parameter	Select an available input parameter from the list. You associate the type of the vRealize Orchestrator workflow with the type of the custom entity definition.
Review Association	Review the details that you specified and click <b>Done</b> to complete the association.

### Example

For example, if you have a custom entity of type `SSH:Host`, you can associate it with the `Add a Root Folder to SSH Host` workflow by selecting the `sshHost` input parameter, which matches the type of the custom entity.

## Dissociate an Action from a Custom Entity Definition

You can remove a vRealize Orchestrator workflow from the list of associated actions.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

### Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 In the card of the selected custom entity, select **Actions > Dissociate Action**.

A new dialog opens.

- 3 Select the workflow you want to remove and click **Dissociate Action**.

The vRealize Orchestrator workflow is no longer associated with the custom entity.

## Publish a Custom Entity

You must publish a custom entity so users from other tenants or service providers can run workflows using the custom entity instances as input parameters.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

### Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 In the card of the selected custom entity, select **Actions > Publish**.

A new dialog opens.

- 3 Choose whether you want to publish the custom entity definition to service providers, all tenants, or only to selected tenants.

- 4 Click **Save** to confirm the change.

The custom entity definition becomes available to the selected parties.

## Delete a Custom Entity

You can delete a custom entity definition if the custom entity is no longer in use, if it was configured incorrectly, or if you want to map the vRealize Orchestrator type to a different custom entity.

### Prerequisites

This operation requires the Custom Entity rights to be included in the predefined user role.

## Procedure

- 1 In the top navigation bar, click **Libraries** and under **Services**, select **Custom Entity Definitions**.

The list of custom entities appears in a card view of 12 items per page, sorted by names alphabetically. Each card shows the name of the custom entity, the vRealize Orchestrator type to which the entity is mapped, the type of the entity, and a description, if available.

- 2 In the card of the selected custom entity, select **Actions > Delete**.

- 3 Confirm the deletion.

The custom entity is removed from the card view.