

VMware Cloud Director Service Provider Admin Portal Guide

Modified on 8 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

Copyright © 2018-2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 VMware Cloud Director™ Service Provider Admin Portal Guide 10**
- 2 Getting Started with VMware Cloud Director Service Provider Admin Portal 11**
 - Overview of VMware Cloud Director Administration 11
 - Log in to VMware Cloud Director Service Provider Admin Portal 15
 - Use the VMware Cloud Director Quick Search 15
 - View Tasks 16
 - Stop a Task in Progress 16
 - View Events 17
 - Set User Preferences 18
 - Length Limits on Names and Descriptions 18
- 3 Managing vSphere Resources 20**
 - Adding vCenter Server and NSX Resources 21
 - Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance 21
 - Discovering and Adopting vApps 25
 - Assign the NSX License Key in vCenter Server 27
 - Register an NSX-T Manager Instance 27
 - Managing NSX Advanced Load Balancing 28
 - Accessing vSphere Components Through VMware Cloud Director Endpoints and Proxies 32
 - Create an Endpoint 33
 - Add a Proxy for Accessing the Underlying vCenter Server Resources 34
 - Manage the Proxy Certificates and CRLs 35
 - Adding Cloud Resources 36
 - Provider Virtual Data Centers 36
 - Create a Provider Virtual Data Center 36
 - External Networks 40
 - Network Pools 43
 - View the vCenter Server Instances 47
 - Modify vCenter Server Settings 48
 - Activate or Deactivate a vCenter Server Instance 49
 - Reconnect a vCenter Server Instance 50
 - Refresh a vCenter Server Instance 50
 - Refresh the Storage Policies of a vCenter Server Instance 50
 - Unregister a vCenter Server Instance 51
 - Modify NSX Manager Settings 51
 - Modify NSX-T Manager Settings 52
 - Delete an NSX-T Manager Instance 52

- Configuring and Managing Multisite Deployments 53
- Multisite Resource Lists 56

4 Managing Provider Virtual Data Centers 57

- Activate or Deactivate a Provider Virtual Data Center 57
- Delete a Provider Virtual Data Center 58
- Edit the General Settings of a Provider Virtual Data Center 58
- Merge Provider Virtual Data Centers 59
- View the Organization Virtual Data Centers of a Provider Virtual Data Center 59
- View the Datastores on a Provider Virtual Data Center 60
- View the External Networks on a Provider Virtual Data Center 61
- Using Kubernetes with VMware Cloud Director 61
 - Creating a vSphere with VMware Tanzu Cluster 65
 - Create a Native Kubernetes Cluster 72
 - Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster 74
- Managing the VM Storage Policies on a Provider Virtual Data Center 75
 - Enabling VM Encryption on Storage Policies of a Provider Virtual Data Center 75
 - Add a VM Storage Policy to a Provider Virtual Data Center 77
 - Activate or Deactivate a VM Storage Policy on a Provider Virtual Data Center 77
 - Delete a VM Storage Policy from a Provider Virtual Data Center 77
 - Modify the Metadata for a VM Storage Policy on a Provider Virtual Data Center 78
 - Enabling the I/O Operations Per Second Setting 78
 - Edit the Provider VDC Storage Policy Settings 80
 - Edit the Entity Types That a Storage Policy Supports 81
- Managing the Resource Pools on a Provider Virtual Data Center 82
 - Add a Resource Pool to a Provider Virtual Data Center 82
 - Activate or Deactivate a Resource Pool on a Provider Virtual Data Center 83
 - Detach a Resource Pool from a Provider Virtual Data Center 84
- Modify the Metadata for a Provider Virtual Data Center 84

5 Managing Organizations 86

- Understanding Leases 86
- Create an Organization 87
- Activate or Deactivate an Organization 87
- Delete an Organization 88
- Configure Catalogs for an Organization 88
- Configure Policies for an Organization 89
- Migrate Tenant Storage 90
- Manage Quotas on the Resource Consumption of an Organization 91

6 Managing Organization Virtual Data Centers 92

Understanding Allocation Models	92
Suggested Use of the Allocation Models	94
Flex Allocation Model	95
Allocation Pool Allocation Model	97
Pay-As-You-Go Allocation Model	98
Reservation Pool Allocation Model	99
Understanding VM Sizing and VM Placement Policies	99
Create a VM Placement Policy within a Provider VDC	103
Create a Global VM Placement Policy	105
Edit a VM Placement Policy	106
Add a VM Placement Policy to an Organization VDC	106
Delete a VM Placement Policy	107
Attributes of VM Sizing Policies	108
Create a VM Sizing Policy	109
Add a VM Sizing Policy to an Organization VDC	110
Edit a VM Sizing Policy	110
Delete a VM Sizing Policy	111
Using Kubernetes with VMware Cloud Director	111
Add an Organization VDC Kubernetes Policy	115
Edit an Organization VDC Kubernetes Policy	117
Create a Tanzu Kubernetes Cluster	117
Create a Native Kubernetes Cluster	119
Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster	120
Create an Organization Virtual Data Center	122
Activate or Deactivate an Organization Virtual Data Center	125
Delete an Organization Virtual Data Center	125
Managing Virtual Data Center Templates	126
Create an Organization Virtual Data Center Template	126
Instantiate a Virtual Data Center from a Template	130
Edit an Organization VDC Template	130
Modify the Name and the Description of an Organization Virtual Data Center	134
Modify the Allocation Model Settings of an Organization Virtual Data Center	134
Modifying the Storage Settings of an Organization Virtual Data Center	135
Enabling VM Encryption on Storage Policies of an Organization Virtual Data Center	135
Modify the VM Provisioning Settings of an Organization Virtual Data Center	136
Add a VM Storage Policy to an Organization Virtual Data Center	136
Change the Default Storage Policy on an Organization Virtual Data Center	137
Edit the Limit of a Storage Policy on an Organization Virtual Data Center	138
Modify the Metadata for a VM Storage Policy on an Organization Virtual Data Center	138
Activate or Deactivate a Storage Policy on an Organization Virtual Data Center	139
Delete a Storage Policy from an Organization Virtual Data Center	139

- Edit the Organization VDC Storage Policy Settings 140
- Edit the Network Settings of an Organization Virtual Data Center 140
- Configuring Cross-Virtual Data Center Networking 142
- Modify the Metadata for an Organization Virtual Data Center 143
- View the Resource Pools of an Organization Virtual Data Center 144
- Managing the Distributed Firewall on an Organization Virtual Data Center 144
 - Activate the Distributed Firewall on an Organization Virtual Data Center 144
 - Add a Distributed Firewall Rule 145
 - Edit a Distributed Firewall Rule 147
 - Custom Grouping Objects 148
 - Working with Security Groups 152
 - Working with Security Tags 155

7 Managing NSX Data Center for vSphere Edge Gateways 160

- Working with NSX Data Center for vSphere Edge Clusters 160
- Add an NSX Data Center for vSphere Edge Gateway 162
- Configuring NSX Data Center for vSphere Edge Gateway Services 164
 - Managing an NSX Data Center for vSphere Edge Gateway Firewall 164
 - Managing NSX Data Center for vSphere Edge Gateway DHCP 168
 - Add a SNAT or a DNAT Rule 173
 - Advanced Routing Configuration 175
 - Load Balancing 184
 - Secure Access Using Virtual Private Networks 197
 - SSL Certificate Management 222
 - Custom Grouping Objects 229
- View the Networks Use and IP Allocations on an Edge Gateway 233
- Editing Edge Gateway Properties 233
 - Activate or Deactivate Distributed Routing on an Edge Gateway 233
 - Modify the External Networks and the Edge Gateway Settings 234
 - Edit the General Settings of an Edge Gateway 234
 - Edit the Default Gateway of an Edge Gateway 235
 - Edit the IP Settings of an Edge Gateway 235
 - Edit the Suballocated IP Pools on an Edge Gateway 236
 - Edit the Rate Limits on an Edge Gateway 236
- Redeploy an Edge Gateway 237
- Delete an Edge Gateway 237
- Statistics and Logs for an Edge Gateway 237
 - View Statistics 237
 - Enable Logging 238
- Enable SSH Command-Line Access to an Edge Gateway 239

8	Managing NSX-T Data Center Edge Gateways	241
	Dedicated External Networks	241
	Add an NSX-T Data Center Edge Gateway	242
	Add an IP Set to an NSX-T Data Center Edge Gateway	243
	Add an NSX-T Data Center Edge Gateway Firewall Rule	243
	Add an SNAT or a DNAT Rule to an NSX-T Edge Gateway	245
	Configure a DNS Forwarder Service on an NSX-T Edge Gateway	248
	Edit the IP Allocations of an NSX-T Edge Gateway	248
	Quick IP Allocation	249
	Create Custom Application Port Profiles	250
	IPsec Policy-Based VPN for NSX-T Data Center Edge Gateways	250
	Configure NSX-T Policy-Based IPsec VPN	251
	Customize the Security Profile of an IPsec VPN Tunnel	252
	Configure Dedicated External Network Services	253
	Manage Route Advertisement	254
	Configure BGP General Settings	254
	Create an IP Prefix List	256
	Add a BGP Neighbor	257
	Managing NSX Advanced Load Balancing on an NSX-T Data Center Edge Gateway	258
	Enable Load Balancer on an NSX-T Data Center Edge Gateway	258
	Assign a Service Engine Group to an NSX-T Data Center Edge Gateway	259
	Edit the Settings of a Service Engine Group	259
	Add a Load Balancer Server Pool	260
	Create a Virtual Service	262
9	Managing Dedicated vCenter Server Instances	265
	Enable the Tenant Access of an Attached vCenter Server	267
	Publish a Dedicated vCenter Server	268
10	Managing System Administrators and Roles	270
	Managing Rights and Roles	270
	Predefined Roles and Their Rights	272
	System Administrator Rights	274
	Rights in Predefined Global Tenant Roles	288
	Managing Rights Bundles	293
	Managing Global Tenant Roles	296
	Managing Provider Roles	300
	Managing Provider Users and Groups	302
	Managing Provider Users	302
	Managing Provider Groups	305

11 Managing System Settings 307

- Modify General System Settings 307
- General System Settings 308
- Activate FIPS Mode on the Cells in the Server Group 310
- Configure the System Email Settings 311
- Change the VMware Cloud Director License 312
- Configure the Catalog Synchronization Settings 313
- Create an Advisory Dashboard 313
- Configuring and Monitoring Blocking Tasks and Notifications 314
 - Configure an AMQP Broker 314
 - Configure Blocking Task Settings 315
 - Monitor Blocked Tasks 316
- Configure Public Addresses 316
- Managing Identity Providers 318
 - Managing LDAP Connections 319
 - Configure Your System to Use a SAML Identity Provider 322
- Managing Certificates 324
 - Import Trusted Certificates 324
 - Import Certificates to the Certificates Library 324
- Managing Plug-Ins 325
 - Upload a Plug-in 326
 - Activate or Deactivate a Plug-in 326
 - Delete a Plug-in 327
 - Publish or Unpublish a Plug-in from an Organization 327
- Customizing the VMware Cloud Director Portals 327
- Configure the Password Policy 329
- Configure vSphere Services 329

12 Monitoring VMware Cloud Director 331

- VMware Cloud Director and Cost Reporting 331
- View Use Information for a Provider Virtual Data Center 332

13 Managing Services 333

- Integrating vRealize Orchestrator with VMware Cloud Director 333
 - Register a vRealize Orchestrator Instance with VMware Cloud Director 334
- Create a Service Category 335
- Edit a Service Category 335
- Import a Service 336
- Search for a Service 336
- Execute a Service 337
- Change a Service Category 338

[Unregister a Service](#) 338

[Publish a Service](#) 339

14 [Managing Defined Entities](#) 340

[Sharing Defined Entities](#) 341

[Managing Custom Entities](#) 343

[Search for a Custom Entity](#) 343

[Edit a Custom Entity Definition](#) 343

[Add a Custom Entity Definition](#) 344

[Custom Entity Instances](#) 345

[Associate an Action to a Custom Entity](#) 345

[Dissociate an Action From a Custom Entity](#) 346

[Publish a Custom Entity](#) 346

[Delete a Custom Entity](#) 347

VMware Cloud Director™ Service Provider Admin Portal Guide

1

The *VMware Cloud Director Service Provider Admin Portal Guide* provides information about how to use the Service Provider Admin Portal. You use the service provider admin portal to manage and monitor organizations, rights, roles, users, and groups in your cloud. You can also create and manage NSX-T backed organization virtual data center networks.

Intended Audience

This guide is intended for service provider administrators who want to use the capabilities provided in the VMware Cloud Director Service Provider Admin Portal.

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

Getting Started with VMware Cloud Director Service Provider Admin Portal

2

The VMware Cloud Director Service Provider Admin Portal is a dedicated interface for service provider administrators.

This chapter includes the following topics:

- [Overview of VMware Cloud Director Administration](#)
- [Log in to VMware Cloud Director Service Provider Admin Portal](#)
- [Use the VMware Cloud Director Quick Search](#)
- [View Tasks](#)
- [Stop a Task in Progress](#)
- [View Events](#)
- [Set User Preferences](#)
- [Length Limits on Names and Descriptions](#)

Overview of VMware Cloud Director Administration

With VMware VMware Cloud Director you can build secure, multi-tenant clouds by pooling virtual infrastructure resources into virtual data centers and exposing them to users through Web-based portals and programmatic interfaces as a fully automated, catalog-based service.

The *VMware Cloud Director Service Provider Admin Portal Guide* provides information about adding resources to the system, creating and provisioning organizations, managing resources and organizations, and monitoring the system.

vSphere and NSX Resources

VMware Cloud Director relies on vSphere resources to provide CPU and memory to run virtual machines. In addition, vSphere datastores provide storage for virtual machine files and other files necessary for virtual machine operations. VMware Cloud Director also uses vSphere distributed switches, vSphere port groups, and NSX Data Center for vSphere to support virtual machine networking.

VMware Cloud Director can also use resources from NSX-T Data Center. For information about registering an NSX-T Manager instance with your cloud, see the *VMware Cloud Director Service Provider Admin Portal Guide* or the *VMware Cloud Director API Programming Guide*.

You can use the underlying vSphere and NSX resources to create cloud resources.

Starting with version 9.7, VMware Cloud Director can act as an HTTP proxy server, with which you can enable organizations to access the underlying vSphere environment.

Cloud Resources

Cloud resources are an abstraction of their underlying vSphere resources. They provide the compute and memory resources for VMware Cloud Director virtual machines and vApps. A vApp is a virtual system that contains one or more individual virtual machines with parameters that define operational details. Cloud resources also provide access to storage and network connectivity.

Cloud resources include provider and organization virtual data centers, external networks, organization virtual data center networks, and network pools.

Before you can add cloud resources to VMware Cloud Director, you must add vSphere resources.

Dedicated vCenter Server Instances and Proxies

A dedicated vCenter Server instance is a cloud resource that encapsulates an entire vCenter Server installation. A dedicated vCenter Server instance includes one or more proxies that are access points to different components of the underlying vSphere environment. The provider can create and enable dedicated vCenter Server instances and proxies. The provider can publish a dedicated vCenter Server instance to tenants.

To create and manage dedicated vCenter Server instances and proxies, you can use the Service Provider Admin Portal or the vCloud OpenAPI. See [Chapter 9 Managing Dedicated vCenter Server Instances](#) and *Getting Started with VMware Cloud Director OpenAPI* at <https://code.vmware.com>.

Provider Virtual Data Centers

A provider virtual data center combines the compute and memory resources of a single vCenter Server resource pool with the storage resources of one or more datastores available to that resource pool.

A provider virtual data center can use network resources from an NSX Manager instance that is associated with the vCenter Server instance or from an NSX-T Manager instance that is registered with the cloud.

You can create multiple provider virtual data centers for users in different geographic locations or business units, or for users with different performance requirements.

Organization Virtual Data Centers

An organization virtual data center provides resources to an organization and is partitioned from a provider virtual data center. Organization virtual data centers provide an environment where virtual systems can be stored, deployed, and operated. They also provide storage for virtual media, such as floppy disks and CD ROMs.

A single organization can have multiple organization virtual data centers.

VMware Cloud Director Networking

VMware Cloud Director supports three types of networks.

- External networks
- Organization virtual data center networks
- vApp networks

Some organization virtual data center networks and all vApp networks are backed by network pools.

External Networks

An external network is a logical, differentiated network based on a vSphere port group. Organization virtual data center networks can connect to external networks to provide Internet connectivity to virtual machines inside a vApp.

Starting with version 9.5, VMware Cloud Director supports IPv6 external networks. An IPv6 external network supports both IPv4 and IPv6 subnets, and an IPv4 external network supports both IPv4 and IPv6 subnets.

By default, only **System Administrators** create and manage external networks.

Organization Virtual Data Center Networks

An organization virtual data center network belongs to 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 in an organization to communicate with each other. To provide external connectivity, you can connect an organization virtual data center network to an external network. You can also create an isolated organization virtual data center network that is internal to the organization.

VMware Cloud Director 9.5 introduces IPv6 support for direct and routed organization virtual data center networks.

Starting with VMware Cloud Director 9.5, **System Administrators** can create isolated virtual data center networks backed by an NSX-T logical switch. **Organization Administrators** can create isolated virtual data center networks backed by network pools.

VMware Cloud Director 9.5 also introduces cross-virtual data center networking by configuring stretched networks in virtual data center groups.

By default, only **System Administrators** can create direct and cross-virtual data center networks. **System Administrators** and **Organization Administrators** can manage organization virtual data center networks, although there are some limits to what an **Organization Administrators** can do.

vApp Networks

A vApp network belongs to a vApp and allows virtual machines in the vApp to communicate with each other. To enable a vApp to communicate with other vApps in the organization, you can connect the vApp network to an organization virtual data center network. If the organization virtual data center network is connected to an external network, the vApp can communicate with vApps from other organizations. vApp networks are backed by network pools.

Most users with access to a vApp can create and manage their own vApp networks. For information about working with networks in a vApp, see *VMware Cloud Director Tenant Portal Guide*.

Network Pools

A network pool is a group of undifferentiated networks that is available for use within an organization virtual data center. A network pool is backed by vSphere network resources such as VLAN IDs or port groups. VMware Cloud Director uses network pools to create NAT-routed and internal organization virtual data center networks and all vApp networks. Network traffic on each network in a pool is isolated at layer 2 from all other networks.

Each organization virtual data center in VMware Cloud Director can have one network pool. Multiple organization virtual data centers can share one network pool. The network pool for an organization virtual data center provides the networks created to satisfy the network quota for an organization virtual data center.

Only **System Administrators** can create and manage network pools.

Organizations

VMware Cloud Director supports multi-tenancy by using 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. **Organization Administrators** tasks are described in *VMware Cloud Director Tenant Portal Guide*.

Users and Groups

An organization can contain an arbitrary number of users and groups. **Organization Administrators** can create users, and import users and groups from a directory service such as LDAP. The **System Administrator** manages the set of rights available to each organization. The **System Administrator** can create and publish global tenant roles to one or more organizations. The **Organization Administrator** can create local roles in their organizations.

Catalogs

Organizations use catalogs to store vApp templates and media files. The members of an organization that have access to a catalog can use the containing vApp templates and media files to create their own vApps. A **System Administrator** can allow an organization to publish a catalog to make it available to other organizations. **Organization Administrators** can then decide which catalog items to provide to their users.

Log in to VMware Cloud Director Service Provider Admin Portal

You can access the VMware Cloud Director Service Provider Admin Portal by using a Web browser.

Prerequisites

You must have the system administrator rights to access the VMware Cloud Director Service Provider Admin Portal.

Procedure

- 1 In a browser, type the Service Provider Admin Portal URL of your VMware Cloud Director site and press Enter.

For example, type `https://vcloud.example.com/provider`.

- 2 Log in with the system administrator user name and password.

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 Service Provider Admin Portal, you can view recent tasks and their status.

You can use the recent tasks view to monitor the status of tasks in your Service Provider Admin Portal. This view can be a good first step for troubleshooting any issues in your environment.

Next to the **Recent Tasks** button, the running and failed tasks appear in blue and red, respectively.

Procedure

- 1 In the lower-left corner, click **Recent Tasks**.
- 2 (Optional) Sort and filter the list of recent tasks.

Results

A lists of recent tasks displays, along with the status of the task, the type, the initiator, and the start and completion time.

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> .

Detail	Description
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**.

Length Limits on Names and Descriptions

Follow these guidelines when entering values in VMware Cloud Director.

String values for the `name` attribute and the `Description` and `ComputerName` elements have length limitations that depend on the object to which they are attached.

Table 2-1. Length Limits on Object Properties

Object	Property	Maximum Length in Characters
Catalog	<code>name</code>	128
Catalog	<code>Description</code>	256
EdgeGateway	<code>name</code>	35

Table 2-1. Length Limits on Object Properties (continued)

Object	Property	Maximum Length in Characters
Media	name	128
Media	Description	256
VApp	name	128
VApp	Description	256
VAppTemplate	name	128
VAppTemplate	Description	256
Vdc	name	128
Vdc	Description	256
Vm	name	128
Vm	ComputerName	15 on Windows, 63 on all other platforms
Vm	Description	256

Managing vSphere Resources

3

VMware Cloud Director derives its resources from an underlying vSphere virtual infrastructure. After you register vSphere resources in VMware Cloud Director, you can allocate these resources for organizations within the vSphere installation to use.

VMware Cloud Director uses one or more vCenter Server environments to back its virtual data centers. Starting with version 9.7, VMware Cloud Director can also use a vCenter Server environment to encapsulate an SDDC with one or more proxies. You can enable tenants to use these proxies as access points to the underlying vSphere environment from VMware Cloud Director with their VMware Cloud Director accounts.

Before you can use a vCenter Server instance in VMware Cloud Director, you must attach this vCenter Server instance.

When you create a provider virtual data center backed by an attached vCenter Server instance, this vCenter Server instance appears as published to a service provider, also called provider scoped. For information about creating a provider virtual data center, see [Create a Provider Virtual Data Center](#).

When you create an SDDC that encapsulates an attached vCenter Server instance, you dedicated the vCenter Server to a tenant. This vCenter Server instance appears as published to a tenant, also called tenant scoped. For information about creating an SDDC, see [Chapter 9 Managing Dedicated vCenter Server Instances](#).

Note By default, with an attached vCenter Server instance, you can create either a provider VDC or a dedicated vCenter Server instance. If you created a provider VDC backed by an vCenter Server instance, you cannot use this vCenter Server instance to create a dedicated vCenter Server instance, and the reverse.

Centralized SSL Management

Starting with version 10.1, VMware Cloud Director is moving to a centralized, tenant-aware storage area for certificate management. This way, VMware Cloud Director centralizes all certificates in one place so that **system administrators** and **organization administrators** can view, audit, and manage all certificates in use by various components in the system. You can use the VMware Cloud Director API to add, update, or remove certificates from the new tenant-aware storage area. See *VMware Cloud Director API Schema Reference*.

When adding or editing a new vCenter Server instance, NSX Manager instance, or NSX-T Manager instance, the VMware Cloud Director UI probes that endpoint for any certificates it is presenting. VMware Cloud Director adds to a centralized certificate storage area any certificate you decide to trust.

This chapter includes the following topics:

- [Adding vCenter Server and NSX Resources](#)
- [Accessing vSphere Components Through VMware Cloud Director Endpoints and Proxies](#)
- [Adding Cloud Resources](#)
- [View the vCenter Server Instances](#)
- [Modify vCenter Server Settings](#)
- [Activate or Deactivate a vCenter Server Instance](#)
- [Reconnect a vCenter Server Instance](#)
- [Refresh a vCenter Server Instance](#)
- [Refresh the Storage Policies of a vCenter Server Instance](#)
- [Unregister a vCenter Server Instance](#)
- [Modify NSX Manager Settings](#)
- [Modify NSX-T Manager Settings](#)
- [Delete an NSX-T Manager Instance](#)
- [Configuring and Managing Multisite Deployments](#)
- [Multisite Resource Lists](#)

Adding vCenter Server and NSX Resources

VMware Cloud Director relies on vSphere resources to provide CPU, memory, and storage to run virtual machines. In addition, starting with version 9.7, VMware Cloud Director can act as an HTTP server between tenants and the underlying vSphere environment.

For information about VMware Cloud Director system requirements and supported versions of vCenter Server and ESXi, see the *VMware Product Interoperability Matrixes* at http://partnerweb.vmware.com/comp_guide/sim/interop_matrix.php.

Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance

You can attach a vCenter Server instance so that its resources become available for use in VMware Cloud Director. You can attach a vCenter Server instance together with its associated NSX Manager instance. For dedicated vCenter Server instances or for those associated with an NSX-T Manager instance, you can attach a vCenter Server instance alone.

VMware Cloud Director can use a vCenter Server instance either with its associated NSX Manager instance or with an NSX-T Manager instance.

If you want VMware Cloud Director to use this vCenter Server instance with its associated NSX Manager instance, you must attach the vCenter Server and NSX Manager instances together.

If you want VMware Cloud Director to use this vCenter Server instance with an NSX-T Manager instance, you must attach the vCenter Server instance alone. After you attach the vCenter Server instance alone, you must [Register an NSX-T Manager Instance](#).

Note After you attach a vCenter Server instance alone, you cannot add its associated NSX Manager instance at a later stage. You can unregister and attach again the vCenter Server instance together with its associated NSX Manager instance.

You can attach a vCenter Server instance to any site from your VMware Cloud Director environment.

You can attach a directly accessible vCenter Server instance or attach a vCenter Server instance that is behind a proxy. By using vCloud OpenAPI, you can use proxy configurations within VMware Cloud Director to create a proxied connection between a VMware Cloud Director instance and the vCenter Server instance added to it. This way, the VMware Cloud Director and vCenter Server instances can exist in different locations or sites.

To attach a vCenter Server instance that is behind a proxy, first, you must declare a proxy configuration. Then, you must attach a vCenter Server instance, and configure VMware Cloud Director to use the proxy configuration when accessing the vCenter Server instance. You can also attach an NSX Data Center for vSphere solution through a proxy. VMware Cloud Director does not support proxy configurations for NSX-T Data Center. You do not need additional SSL configurations or an additional proxy configuration for the Platform Services Controller the vCenter Server instance is registered with.

Prerequisites

- If you configured VMware Cloud Director to verify vCenter and vSphere SSO certificates, verify that you uploaded the vCenter Server certificates to VMware Cloud Director. For information about the general system settings, see [Modify General System Settings](#).
- If you configured VMware Cloud Director to verify NSX Manager certificates, verify that you uploaded the NSX Manager certificates to VMware Cloud Director. For information about the general system settings, see [Modify General System Settings](#).

Procedure

1 [Add the vCenter Server Instance](#)

To add a vCenter Server instance, you enter the vCenter Server access details.

2 [\(Optional\) Add the Associated NSX Manager Instance](#)

If you want VMware Cloud Director to use this vCenter Server instance with its associated NSX Manager instance, you must add NSX Manager access details.

Add the vCenter Server Instance

To add a vCenter Server instance, you enter the vCenter Server access details.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left pane, click **vCenter Server Instances** and click **Add**.
- 3 If you have a multisite VMware Cloud Director deployment, from the **Site** drop-down menu, select the site to which you want to add this vCenter Server instance, and click **Next**.
- 4 Enter a name and, optionally, a description for the vCenter Server instance in VMware Cloud Director.
- 5 Enter the URL of the vCenter Server instance.
If the default port is used, you can skip the port number. If a custom port is used, include the port number
For example, `https://FQDN_or_IP_address:<custom_port_number>`.
- 6 Enter the user name and password of the vCenter Server **administrator** account.
- 7 (Optional) To deactivate the vCenter Server instance after the registration, turn off the **Enabled** toggle.
- 8 Configure the URL of the vCenter Server Web Client.

Option	Description
Use vSphere Services to provide URL	To use this option, you must use the vCloud API to configure VMware Cloud Director to use the vSphere Lookup Service.
vSphere Web Client URL	To use this option, you must enter the URL of the vSphere Web Client. For example, <code>https://example.vmware.com/vsphere-client</code> .

- 9 Click **Next**.
- 10 If the endpoint does not have a trusted certificate, on the **Trust Certificate** window confirm if you trust the endpoint.

In a multisite environment, if you are logged in to a vCloud Director 10.0 site or trying to register a vCenter Server instance to a vCloud Director 10.0 site, VMware Cloud Director will not add the endpoint to the centralized certificate storage area.

- To add the endpoint to the centralized certificate storage area and continue, click **Trust**.
- If you do not trust this endpoint, click **Cancel** and repeat [Step 5](#) to [Step 9](#) with a trusted endpoint.

- 11 (Optional) Skip adding the NSX Manager instance that is associated with the vCenter Server instance by turning off the **Configure Settings** toggle and click **Next**.

If you want VMware Cloud Director to use this vCenter Server instance with an NSX-T Manager instance, you must add the vCenter Server instance alone.

Note You cannot add the associated NSX Manager instance at a later stage. You can unregister and attach again the vCenter Server instance together with its associated NSX Manager instance.

- 12 If you want to add a tenant dedicated vCenter Server that will not be used as a provider VDC, turn on the **Enable tenant access** toggle.

After you add the vCenter Server instance to VMware Cloud Director, the tenant-related information appears in the details view of the instance.

- 13 If you want VMware Cloud Director to generate default proxies for the vCenter Server instance and SSO services, turn on the **Generate proxies** toggle.

After you add the vCenter Server instance to VMware Cloud Director, the proxies appear in the **Proxies** tab under **vSphere Resources**.

- 14 On the **Ready to Complete** page, review the registration details and click **Finish**.

(Optional) Add the Associated NSX Manager Instance

If you want VMware Cloud Director to use this vCenter Server instance with its associated NSX Manager instance, you must add NSX Manager access details.

Procedure

- 1 On the **NSX-V Manager** page, leave the **Configure Settings** toggle turned on.
- 2 Enter the URL of the NSX Manager instance.

If the default port is used, you can skip the port number. If a custom port is used, include the port number

For example, `https://FQDN_or_IP_address:<custom_port_number>`.

- 3 Enter the user name and password of the NSX **administrator** account.

- 4 (Optional) To enable cross-virtual data center networking for the virtual data centers backed by this vCenter Server instance, turn on the **Cross-VDC networking** toggle, and enter the control VM deployment properties and a name for the network provider scope.

The control VM deployment properties are used for deploying an appliance on the NSX Manager instance for cross-virtual data center networking components like a universal router.

Option	Description
Network Provider Scope	Corresponds to the network fault domain in the network topologies of the data center groups. For example, boston-fault1 . For information about managing cross-virtual data center groups, see the <i>VMware Cloud Director Tenant Portal Guide</i> .
Resource Pool Path	The hierarchical path to a specific resource pool in the vCenter Server instance, starting from the cluster, <i>Cluster/Resource_Pool_Parent/Target_Resource</i> . For example, TestbedCluster1/mgmt-rp . As an alternative, you can enter the Managed Object Reference ID of the resource pool. For example, resgroup-1476 .
Datastore Name	The name of the datastore to host the appliance files. For example, shared-disk-1 .
Management Interface	The name of the network in vCenter Server or port group used for the HA DLR management interface. For example, TestbedPG1 .

- 5 Click **Next**.
- 6 If the endpoint does not have a trusted certificate, on the **Trust Certificate** window confirm if you trust the endpoint.
- To add the endpoint to the centralized certificate storage area and continue, click **Trust**.
 - If you do not trust this endpoint, click **Cancel** and repeat **Step 2** to **Step 4** with a trusted endpoint.
- 7 Activate or deactivate the access configuration settings.
- 8 On the **Ready to Complete** page, review the registration details and click **Finish**.

What to do next

- [Assign the NSX License Key in vCenter Server](#).
- [Create a Provider Virtual Data Center](#).

Discovering and Adopting vApps

In the default configuration, an organization VDC discovers VMs that are created in any vCenter Server resource pool that backs the VDC. The system constructs a simplified vApp, owned by the system administrator, to contain each discovered virtual machine (VM). After the system administrator grants you access to a discovered vApp, you can reference the VM in it when you compose or recompose a vApp, or modify the vApp to adopt it and import it.

Discovered vApps contain exactly one VM, and are subject to several constraints that do not apply to vApps created in VMware Cloud Director. Whether or not you adopt them, they can be useful as a source of VMs to use when composing or recomposing a vApp.

Each discovered vApp is given a name that is derived from the name of the vCenter VM that it contains and a prefix specified by your organization administrator.

If you want to discover additional vApps, a system administrator can use the VMware Cloud Director API to create organization VDCs that adopt specified resource pools available from a Provider VDC. vCenter VMs in these adopted resource pools appear in the new VDC as discovered vApps, and are candidates for adoption.

Note Virtual machines with IDE hard drives are discovered only if they are in powered off state.

If one or more vCenter VMs are not discovered by VMware Cloud Director, you can investigate the possible reasons by debugging the vCenter Server VM Discovery. For more information, see the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

Activating VM Discovery

VM discovery is activated by default. To deactivate VM discovery, a system administrator must deselect the **VM discovery enabled** check box on the **System Settings > General** tab. An organization administrator can use the VMware Cloud Director API to deactivate VM discovery for individual VDCs, or for all VDCs in an organization.

Using a VM from a Discovered vApp

After the system administrator grants you access to a discovered vApp, you can use its VM in the same ways you can use a VM that any other vApp or vApp template contains. For example, you can specify it when you build a new vApp. You can also clone a discovered vApp or modify its name, description, or lease settings without triggering the adoption process.

Adopting a Discovered vApp

You can adopt a discovered vApp by changing its vApp network or adding a VM to this vApp. After you adopted a discovered vApp, the system imports it and treats it as though it was created in VMware Cloud Director. When an adopted vApp is retrieved with a vCloud API request, it includes an element named `autoNature`. This element has a value of `false` if the discovered vApp was adopted or was created in VMware Cloud Director. You cannot revert an adopted vApp to a discovered vApp.

If you delete or move the VM that a discovered vApp contains, the system also removes the containing vApp. This behavior does not apply to adopted vApps.

The vApp created to contain a discovered vCenter VM is similar to the one created when you manually import a VM as a vApp, but it is simplified in ways that might require you to modify it before you can deploy it in your VDC. For example, you might have to edit its networking and storage properties, and make other adjustments specific to the needs of your organization.

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

Assign the NSX License Key in vCenter Server

If you attached a vCenter Server instance together with its associated NSX Manager instance, you must use the vSphere Client to assign a license key for the NSX Manager instance that supports VMware Cloud Director networking.

Prerequisites

This operation is restricted to system administrators.

Procedure

- 1 From a vSphere Client that is connected to the vCenter Server system, select **Home > Licensing**.
- 2 For the report view, select **Asset**.
- 3 Right-click the NSX Manager asset and select **Change license key**.
- 4 Select **Assign a new license key** and click **Enter Key**.
- 5 Enter the license key, enter an optional label for the key, and click **OK**.

Use the NSX Manager license key you received when you purchased VMware Cloud Director. You can use this license key in multiple vCenter Server instances.

- 6 Click **OK**.

Register an NSX-T Manager Instance

You can register an NSX-T Manager instance with VMware Cloud Director, so that VMware Cloud Director can use its network resources. A provider virtual data center can use network resources either from NSX Data Center for vSphere or from NSX-T Data Center.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left pane, click **NSX-T Managers** and click **Add**.
- 3 If you have a multisite VMware Cloud Director deployment, from the **Site** drop-down menu, select the site to which you want to add this NSX-T Manager instance, and click **Next**.

- 4 Enter a name and, optionally, a description for the NSX-T Manager instance in VMware Cloud Director.
- 5 Enter the URL of the NSX-T Manager instance.
For example, `https://FQDN_or_IP_address`.
- 6 Enter the user name and password of the NSX-T Manager **administrator** account.
- 7 Click **Save**.

What to do next

For information about creating a provider virtual data center backed by NSX-T Data Center, see *VMware Cloud Director API Programming Guide* at <https://code.vmware.com>.

Managing NSX Advanced Load Balancing

Starting with version 10.2, VMware Cloud Director provides load balancing services by leveraging the capabilities of VMware NSX Advanced Load Balancer.

As a **system administrator**, you can enable and configure access to load balancing services for virtual data centers backed by NSX-T Data Center.

Load balancing services are associated with NSX-T Data Center edge gateways, which can be scoped either to an organization VDC backed by NSX-T Data Center or to a data center group with NSX-T Data Center network provider type.

After you deploy and configure NSX Advanced Load Balancer to use with your NSX-T Data Center deployment, you register Controllers with VMware Cloud Director.

For information on how to configure NSX Advanced Load Balancer with NSX-T, see [Avi Integration with NSX-T](#).

For information about how to deploy NSX Advanced Load Balancer with VMware Cloud Director, see [Deploying NSX Advanced Load Balancer with VMware Cloud Director](#).

To use the virtual infrastructure provided by NSX Advanced Load Balancer, register your NSX-T Cloud instances with VMware Cloud Director. Controllers serve as a central control plane for load balancing services. After you register your controllers, you can manage them directly from VMware Cloud Director.

The load balancing compute infrastructure provided by NSX Advanced Load Balancer is organized into service engine groups. You can assign more than one service engine group to an NSX-T Data Center edge gateway in VMware Cloud Director. All service engine groups that are assigned to a single edge gateway use the same network.

A service engine group has a unique set of compute characteristics that you define upon creation.

After a **system administrator** assigns a service engine group to an edge gateway, an **organization administrator** can create and configure virtual services that run in a specific service engine group.

Register a Controller Instance

To integrate VMware Cloud Director with your NSX Advanced Load Balancer deployment, you register Controller instances with your VMware Cloud Director instance.

Controller instances serve as a central control plane for the load-balancing services provided by NSX Advanced Load Balancer.

Prerequisites

Install and configure NSX Advanced Load Balancer with your NSX-T Data Center instance.

For information on how to configure NSX Advanced Load Balancer with NSX-T, see [Avi Integration with NSX-T](#).

Note The FQDN or IP address that you use to register NSX-T Manager with NSX Advanced Load Balancer must match the FQDN or IP address of the NSX-T Manager instance that you used to register NSX-T Data Center with VMware Cloud Director.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 Click **NSX-ALB** and then click **Controllers**.
- 3 To add a controller, click **Add**.
- 4 If you are using a multisite deployment, from the drop-down menu, select a site in which to register the Controller.
- 5 Register the Controller instance.
 - a Enter a meaningful name and, optionally, a description for the Controller instance.
 - b Enter the URL of the Controller.
For example, `https://FQDN-or-IP-address`.
 - c Enter the user name and password for the Controller.
 - d Click **Save**.

Results

The Controller instance appears in the list as enabled.

What to do next

[Register an NSX-T Cloud](#).

Register an NSX-T Cloud

To use the virtual infrastructure provided by NSX Advanced Load Balancer, register your NSX-T Cloud instances with VMware Cloud Director.

An NSX-T Cloud is a service provider-level construct that consists of an NSX-T Manager and an NSX-T Data Center transport zone.

NSX-T Manager provides a system view and is the management component of NSX-T Data Center. An NSX-T Data Center transport zone dictates which hosts and virtual machines can participate in the use of a particular network.

If there are multiple transport zones managed by the same NSX-T Manager, then a separate NSX-T Cloud encapsulates each pair of NSX-T Manager and NSX-T Data Center transport zone instances.

An NSX-T Cloud has a one-to-one relationship with a network pool backed by an NSX-T Data Center transport zone.

Prerequisites

[Register a Controller Instance.](#)

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 Click **NSX-ALB** and then click **NSX-T Clouds**.
- 3 To add an NSX-T cloud, click **Add**.
- 4 From the drop-down menu, select a Controller instance for which to create the NSX-T Cloud.
- 5 Enter a name and, optionally, a description for the NSX-T Cloud.
- 6 Select an available Cloud from the list.
- 7 To import the cloud, click **Add**.

Results

The imported cloud appears in the list of available NSX-T Clouds.

What to do next

[Import a Service Engine Group.](#)

Import a Service Engine Group

To provide virtual service management capabilities to your tenants, import service engine groups to your VMware Cloud Director deployment.

A service engine group is an isolation domain that also defines shared service engine properties, such as size, network access, and failover.

Resources in a service engine group can be used for different virtual services, depending on your tenant needs. These resources cannot be shared between different service engine groups.

You can manage and update service engine groups by using NSX Advanced Load Balancer. After you update a service engine group in NSX Advanced Load Balancer, you must sync it to update its settings in the VMware Cloud Director UI.

Only an imported service engine group can be assigned to an edge gateway.

To import a service engine group, associate it with an NSX-T Cloud that is already registered with your VMware Cloud Director instance.

Prerequisites

- [Register a Controller Instance.](#)
- [Register an NSX-T Cloud.](#)

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 Click **NSX-ALB** and then click **Service Engine Groups**.
- 3 To import a service engine group, click **Add**.
- 4 From the drop-down menu, select an NSX-T Cloud.
- 5 Select a reservation model.
 - To assign the service engine group to a single edge gateway, select **Dedicated**.
 - To share the service engine group between several edge gateways, select **Shared**.
- 6 Enter a name and, optionally, a description, for the service engine group.
- 7 Select a service engine group instance.
- 8 Click **Add**.

What to do next

Enable load balancing on the edge gateway and assign the service engine group to the edge gateway. See [Managing NSX Advanced Load Balancing on an NSX-T Data Center Edge Gateway](#).

Sync a Service Engine Group

To update the settings of an imported service engine group, you must sync it with NSX Advanced Load Balancer.

You can manage and update service engine groups by using NSX Advanced Load Balancer. After you update a service engine group in NSX Advanced Load Balancer, you must sync it to update its settings in the VMware Cloud Director UI.

Syncing a service engine group updates the local record of the group's high availability mode and the maximum number of virtual services that the service engine group supports.

Important 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

[Import a Service Engine Group.](#)

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 Select **NSX-ALB** and then click **Service Engine Groups**.
- 3 Select a service engine group and click **Sync**.

Results

The settings of the service engine group are updated.

Accessing vSphere Components Through VMware Cloud Director Endpoints and Proxies

You can use VMware Cloud Director endpoints to access the underlying vSphere environment. When endpoints are connected to proxies, VMware Cloud Director acts as an HTTP proxy server.

Endpoints

A VMware Cloud Director endpoint is an access point to a data center component, for example, a vCenter Server instance, an ESXi host, or an NSX Manager instance. Users can log in to the UI or API of proxied or non-proxied components by using their VMware Cloud Director accounts.

Creating a dedicated vCenter Server instance also creates a default endpoint for it. While attaching the vCenter Server instance, you can also create a proxy. However, the default endpoint is not connected to any proxy by default. You must edit the default endpoint or create a new one to connect it to a proxy.

You can create, edit, and delete endpoints from the **Endpoints** tab of a dedicated vCenter Server instance. See [Create an Endpoint](#).

Proxies

The VMware Cloud Director provided proxies are different from the proxy configurations within VMware Cloud Director. Unlike VMware Cloud Director provided proxies that are scoped to a tenant, proxy configurations within VMware Cloud Director are on the provider level and there is no tenancy.

By activating and deactivating a VMware Cloud Director provided proxy, you can allow and stop the tenant access through that proxy.

You can create a proxy either when you attach a vCenter Server instance to VMware Cloud Director or later. If you create a proxy while attaching a vCenter Server and activating the tenant access, you must manually connect the proxy to the default endpoint.

If the vCenter Server instance uses an external Platform Services Controller, VMware Cloud Director creates a proxy for the Platform Services Controller as well. With parent and child proxies, you can hide certain proxies from the tenants or you can activate and deactivate groups of child proxies through their parent proxies. For information on creating a proxy after you add a vCenter Server instance to VMware Cloud Director, see [Add a Proxy for Accessing the Underlying vCenter Server Resources](#).

You can edit, activate, deactivate, and delete proxies from the **Proxies** tab under **Infrastructure Resources**.

Note When you add a proxy to a vCenter Server instance, you must upload the certificate and the thumbprint, so that tenants can retrieve the certificate and the thumbprint if the proxied component uses self-signed certificates.

To view and manage certificates and certificate revocation lists (CRLs), see [Manage the Proxy Certificates and CRLs](#).

Create an Endpoint

You can create endpoints that administrators and tenants can use to access the underlying vSphere environment.

Endpoints must be attached to dedicated vCenter Server instances and are visible to the tenants from the **Actions** menu of the dedicated vCenter Server instances. If you enable the tenant access when you add a vCenter Server instance to VMware Cloud Director, VMware Cloud Director creates a default endpoint with the vCenter Server instance URL as a target URL. If you create additional endpoints, you can change the default one.

Endpoints can serve as links between dedicated vCenter Server instances and proxies. Endpoints can have a connection to one proxy or they might not have a proxy connection. If an endpoint is connected to a proxy, the target of the endpoint is the target URL, not the UI URL of the connected proxy.

Prerequisites

Verify that the vCenter Server instance for which you want to create endpoints has enabled tenant access. See [Enable the Tenant Access of an Attached vCenter Server](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.
- 3 Select a vCenter Server instance.
- 4 On the page with detailed vCenter Server information, click the **Endpoints** tab and click **New**.
- 5 Enter a name and a target URL for the endpoint.
- 6 (Optional) Make this endpoint the default endpoint for this vCenter Server instance.
- 7 (Optional) Make a connection to a proxy.
- 8 Click **Save**.

What to do next

- Edit the endpoint settings.
- Delete an endpoint. If you want to delete the default endpoint, you must select another one as the default.

Add a Proxy for Accessing the Underlying vCenter Server Resources

If you want VMware Cloud Director to act as an HTTP proxy server for vCenter Server instances and their components, you can create a proxy. You can create proxies for dedicated vCenter Server instances and the vCenter Server instances that do not have a set purpose.

If you want to generate automatically a vCenter Server proxy with retrieved certificates and thumbprint, you can do so from the **vCenter Server Instances** grid or the vCenter Server details view. If the vCenter Server is with an external Platform Services Controller, this option also creates a proxy for the SSO endpoint.

This procedure describes how to create manually a proxy for a vCenter Server instance, or create a proxy for an ESXi host, external Platform Services Controller instance, or NSX Manager instance.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.
- 3 Select a vCenter Server instance.
- 4 On the page with detailed vCenter Server information, click the **Proxies** tab and click **New**.
- 5 Enter a name for the proxy.

- 6 Select the type of the proxy, depending on the component that you want VMware Cloud Director to be a proxy for.

You cannot edit this setting after the creation of the proxy.

You can create only one vCenter Server proxy. If there is an existing vCenter Server proxy and you want to create a new proxy, the **Type** drop-down menu does not include a vCenter Server option.

- If you want to create a vCenter Server proxy, select **vCenter** from the **Type** drop-down menu and continue to [Step 10](#).
- If you want to create a proxy for an ESXi host, NSX Manager, or SSO, make your selection from the drop-down menu and continue to [Step 7](#).

- 7 Enter a name, target host, and the UI URL of the new proxy.

The target host is the host name or IP address of the component that you want VMware Cloud Director to be a proxy for. The UI URL of the new proxy is the URL to which the VMware Cloud Director UI directs to when the tenant opens the proxy.

- 8 If you want the proxy to be visible to the tenants, toggle on the **Tenant visible** option.
- 9 (Optional) Click **Select a parent proxy** and select a proxy from the list.

- 10 Click **Save**.

What to do next

[Manage the Proxy Certificates and CRLs.](#)

Manage the Proxy Certificates and CRLs

You can view, download, and upload the proxy certificates and certificate revocation lists (CRLs).

Prerequisites

Verify that you have VMware Cloud Director provided proxies for at least one vCenter Server instance. See [Accessing vSphere Components Through VMware Cloud Director Endpoints and Proxies](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, click **Proxies**, and select a proxy.
- 3 Click **Manage Certificate**.
- 4 Upload or download the certificate and CRL.
- 5 Click **Save**.

Adding Cloud Resources

Cloud resources are an abstraction of their underlying vSphere resources and provide the compute and memory resources for VMware Cloud Director virtual machines and vApps, and access to storage and network connectivity.

Cloud resources include provider and organization virtual data centers, external networks, organization virtual data center networks, and network pools. Before you can add cloud resources to VMware Cloud Director, you must add vSphere resources.

For information about organization virtual data centers, see [Chapter 6 Managing Organization Virtual Data Centers](#).

For information about organization virtual data center networks, see the *Managing Organization Virtual Data Center Networks* chapter in the *VMware Cloud Director Tenant Portal Guide*.

VMware Cloud Director 9.7 introduces the SDDC or dedicated vCenter Server instance as a cloud resource that encapsulates an entire vCenter Server installation. The provider can create and enable a dedicated vCenter Server, publish it to tenants, and create and enable proxies to different components of the underlying vSphere environment. To create, publish to tenants, and manage dedicated vCenter Server instances and proxies, you can use the Service Provider Admin Portal or vCloud OpenAPI. See [Chapter 9 Managing Dedicated vCenter Server Instances or Getting Started with VMware Cloud Director OpenAPI](#) at <https://code.vmware.com>.

Provider Virtual Data Centers

A provider virtual data center (VDC) combines the compute and memory resources of vCenter Server resource pools with the storage resources of one or more storage policies from a single vCenter Server instance. For network resources, a provider VDC can use either NSX Data Center for vSphere or NSX-T Data Center.

- You can create and manage a provider VDC backed by an attached vCenter Server instance and its associated NSX Manager instance by using the Service Provider Admin Portal or the vCloud API.
- You can create and manage a provider VDC backed by an attached vCenter Server instance and an NSX-T Manager instance by using the Service Provider Admin Portal or the vCloud API.

A typical VMware Cloud Director system includes multiple provider VDCs configured to meet various service level requirements. Each provider VDC has a primary resource pool. You can add and remove non-primary resource pools from the backing vCenter Server instance. You cannot remove the primary resource pool.

Create a Provider Virtual Data Center

To make vSphere compute, memory, and storage resources available to VMware Cloud Director, you create a provider virtual data center (VDC).

Before an organization can begin deploying VMs or creating catalogs, the **system administrator** must create a provider VDC and the organization VDCs that consume its resources. The relationship of provider VDCs to the organization VDCs they support is an administrative decision. The decision can be based on the scope of your service offerings, the capacity and geographical distribution of your vSphere infrastructure, and similar considerations. Because a provider VDC constrains the vSphere capacity and services available to tenants, **system administrators** commonly create provider VDCs that furnish different classes of service, as measured by performance, capacity, and features. Tenants can then be provisioned with organization VDCs that deliver specific classes of service defined by the configuration of the backing provider VDC.

Before you create a provider VDC, consider the set of vSphere capabilities that you plan to offer your tenants. Some of these capabilities can be implemented in the primary resource pool of the provider VDC. Others might require you to create additional resource pools based on specially configured vSphere clusters and add them to the VDC as described in [Add a Resource Pool to a Provider Virtual Data Center](#).

The range of ESXi releases installed on hosts in the cluster backing a resource pool determines the set of guest operating systems and virtual hardware versions available to VMs deployed in organization VDCs backed by the provider VDC.

Prerequisites

- Log in to the Service Provider Admin Portal as a **system administrator**.
- Verify that you created the target primary resource pool with available capacity in a cluster configured to use automated DRS. You can use a resource pool for only one provider VDC. To create a resource pool, you can use the vSphere Client.

If you plan to use a resource pool that is part of a cluster that uses vSphere High Availability (HA), verify that you are familiar with how vSphere HA calculates the slot size. For information about slot sizes and customizing vSphere HA behavior, see the *vSphere Availability* documentation.

- If you want to use vSphere with VMware Tanzu in VMware Cloud Director, verify that you have available a vCenter Server 7.0 or later instance with a configured Supervisor Cluster. See the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation.
- If you use NSX Data Center for vSphere for the network resources of the provider VDC:
 - Verify that the vCenter Server instance that contains the target primary resource pool is attached and has a NSX Data Center for vSphere license key.
 - Set up the VXLAN infrastructure in NSX Manager. See the relevant *NSX Administration Guide*.

If you want to use a custom VXLAN network pool in this provider VDC instead of the default VXLAN network pool, create that network pool now. See [Create a Network Pool Backed by an NSX Data Center for vSphere Transport Zone](#).

- If you use NSX-T Data Center for the network resources of the provider VDC:
 - [Add an External Network That Is Backed by an NSX-T Data Center Tier-0 Gateway](#)
 - [Create a Network Pool Backed by an NSX-T Data Center Transport Zone](#)

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**.
- 3 Click **New**.
- 4 If you have a multisite VMware Cloud Director deployment, from the **Site** drop-down menu, select the site to which you want to add this provider VDC instance, and click **Next**.
- 5 Enter a name and, optionally, a description for the provider VDC.

You can use these text boxes to indicate the vSphere features available to organization VDCs backed by this provider VDC, for example, **vSphere HA** or **Storage policies with IOPS support**.

- 6 (Optional) To deactivate the provider VDC upon creation, turn off the **State** toggle.
You cannot use the compute and storage resources of a deactivated VDC for the creation of organization VDCs.
- 7 Click **Next**.
- 8 To provide resource pools for the provider VDC, select a vCenter Server instance, and click **Next**.

This page lists vCenter Server instances registered to VMware Cloud Director. Clicking a vCenter Server instance shows its available resource pools.

If you want to use vSphere with VMware Tanzu in VMware Cloud Director, you must select a vCenter Server 7.0 or later instance with a configured Supervisor Cluster.

- 9 Select a resource pool to serve as the primary resource pool for this provider VDC.
You can use one resource pool for one provider VDC. When you add a resource pool to a provider VDC, this resource pool and its parent chain become unavailable for selection for other provider VDCs.
If you want to use vSphere with VMware Tanzu, select a Supervisor Cluster. VMware Cloud Director displays a Kubernetes icon next to resource pools backed by a Supervisor Cluster.
- 10 If you select a resource pool or cluster that is backed by a Supervisor Cluster, to establish a trust relationship with the Kubernetes control plane, you must trust the Kubernetes control plane certificate.

- 11 Select the highest virtual hardware version you want the provider VDC to support, and click **Next**.

The system determines the highest virtual hardware version supported by all hosts in the cluster that backs the resource pool and offers it as the default in the **Highest supported hardware version** drop-down menu. You can use this default or select a lower hardware version from the menu. The version you specify becomes the highest virtual hardware version available to a VM deployed in an organization VDC backed by this provider VDC. If you select a lower virtual hardware version, some guest operating systems might not be supported for use by those VMs. Once you create the provider VDC with the selected hardware version, you can only upgrade the version, you cannot downgrade it.

Note The available hardware version for the provider VDC depends on the highest available version of the ESXi host in the target cluster. If the highest supported hardware version of the ESXi host is not available for selection, verify in the vSphere Client that the default compatibility for virtual machine creation on the data center is set to **Use datacenter setting and host version**. You can also set the default compatibility setting to the highest hardware version you want for the cluster.

VMware Cloud Director 9.7 and later support the highest hardware version that the backing vSphere infrastructure supports. Starting with VMware Cloud Director 10.2.2, you can set the hardware version without manually configuring the default hardware version in the vCenter Server instance.

- 12 Select one or more storage policies for the provider VDC, and click **Next**.

All vSphere storage policies supported by the resource pool you selected are listed.

- 13 Configure the network pool for this provider VDC.

Every provider VDC must have a network pool. You can have the system create one for you with a default scope, or you can use a custom VXLAN based on a specific NSX Data Center for vSphere or a Geneve pool based on a NSX-T Data Center transport zone.

Note If you want to use vSphere with VMware Tanzu in VMware Cloud Director, you must select the **NSX-T Manager and Geneve Network pool** option.

Option	Description
Create a default VXLAN Network Pool	The system creates a VXLAN pool for this provider VDC.
Select VXLAN Network Pool from list	You select a network pool from a list so that you use a custom VXLAN pool based on a specific NSX transport zone.
Select NSX-T Manager and Geneve Network pool	You select a network pool from a list so that you use a custom VXLAN pool backed by an NSX-T Data Center transport zone.

- 14 Review your choices and click **Finish** to create the provider VDC.

What to do next

You can add secondary resource pools that enable the provider VDC to provide specialized capabilities such as Edge clusters, affinity groups, and hosts with special configurations that some organizations might require. See [Add a Resource Pool to a Provider Virtual Data Center](#).

External Networks

A VMware Cloud Director external network provides an uplink interface that connects networks and virtual machines in the system to a network outside of the system, such as a VPN, a corporate intranet, or the public Internet. Only a **system administrator** can create an external network.

If you have more than one vCenter Server instance registered to the system, you can create multiple external networks, each backed either by a vSphere network or a tier-0 logical router.

VMware Cloud Director supports IPv4 and IPv6 external networks.

Note The range of IP addresses that you define when you create the external network are allocated either to an edge gateway or to the virtual machines that are directly connected to the network. Because of this, the IP addresses must not be used outside of VMware Cloud Director.

External Networks Backed by vSphere Networks

External networks can be backed either by a single vSphere network, or by multiple vSphere networks.

- External networks backed by a single vSphere instance.

To provide each consumer of the external network with a non-overlapping set of IP addresses on the vSphere network, the **system administrator** must configure the IP ranges on the underlying VLAN manually.

- External networks backed by multiple vSphere networks.

An external network can be backed by multiple vSphere networks. This approach can simplify the IP address management in VMware Cloud Director. You can modify the properties of an external network to change its network backings.

External networks backed by multiple vSphere networks have several constraints.

- A network can have at most one backing vSphere network on each VMware Cloud Director instance registered to the system.
- All backing network switches must be of the same type, either vSphere Distributed Switch or standard switch.

External Networks Backed by a Tier-0 Logical Router

An external network can be backed by an NSX-T Data Center tier-0 logical router.

You can also create an external network that is backed by a VRF-lite tier-0 gateway in NSX-T Data Center.

A virtual routing and forwarding (VRF) gateway is created from a parent tier-0 gateway. It has its own routing tables.

Multiple VRF gateways can exist within the same tier-0 gateway at the same time. Because of that, creating a VRF-backed external network makes possible the creation of a fully routed network topology in a VDC by scaling out a tier-0 gateway in NSX-T Data Center.

For information about VRF gateways, see *NSX-T Data Center Administration Guide*.

Add an External Network That Is Backed by vSphere Resources

By adding an external network, you can register vSphere network resources for VMware Cloud Director to use. You can create organization VDC networks that connect to an external network.

You can add an IPv4 or IPv6 external network. An IPv6 external network supports both IPv4 and IPv6 subnets, and an IPv4 external network supports both IPv4 and IPv6 subnets.

Prerequisites

Verify that a vSphere port group is available with or without VLAN trunking. Elastic port groups with static port binding ensure optimal performance.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left pane, click **External Networks** and click **New**.
- 3 Select **vSphere Resources**, select the type of port groups to back the network, and click **Next**.
- 4 Enter a name and, optionally, a description for the new external network.
- 5 Select the port groups to back the external network and click **Next**.
- 6 Configure at least one subnet and click **Next**.
 - a To add a subnet, click **Add**.
 - b Enter the network Classless Inter-Domain Routing (CIDR) settings.
Use the format *network_gateway_IP_address/subnet_prefix_length*, for example, **192.167.1.1/24**.
 - c (Optional) Enter the DNS settings.
 - d Configure a static IP pool by adding at least one IP range or IP address.
 - e Click **OK**.
 - f (Optional) To add another subnet, repeat this step.
- 7 Review the network settings and click **Finish**.

What to do next

You can create an organization VDC network that connects to the external network.

Add an External Network That Is Backed by an NSX-T Data Center Tier-0 Gateway

To register NSX-T Data Center network resources for VMware Cloud Director to use, add an external network backed by a tier-0 gateway.

Prerequisites

To create an external network that is backed by an NSX-T Data Center tier-0 gateway, first you must create a tier-0 gateway. You can create the tier-0 gateway in the NSX-T Manager UI or by using the NSX Policy API.

If you want to create an external network that is backed by a VRF gateway in NSX-T Data Center, you must also create a VRF gateway that is linked to the tier-0 gateway.

- Create a tier-0 gateway in the NSX-T Manager UI .
 - a Log in with administrative privileges to the NSX-T Manager instance.
 - b Click **Networking**, click **Tier-0 Gateways**, and click **AddGateway > Tier-0** .
 - c Enter a name for the Tier-0 router.
 - d Select a High Availability mode.

Note By default, the active-active mode is used. In the active-active mode, the traffic is load balanced across all members. In active-standby mode, an elected active member processes the traffic. If the active member fails, a new member becomes active.

 - e Select an existing NSX-T Edge cluster from the drop-down menu to back this tier-0 logical router, and click **Save**.
- If you want to create an external network that is backed by a VRF gateway in NSX-T Data Center, create a VRF gateway that is linked to the tier-0 gateway.
 - a Log in with administrative privileges to the NSX-T Manager instance.
 - b Click **Networking**, click **Tier-0 Gateways**, and click **Add Gateway > VRF**.
 - c Enter a name for the VRF gateway.
 - d Select the tier-0 gateway to which to connect the VRF gateway.
 - e Click **Save**.

Procedure

- 1 Log in to the VMware Cloud Director Service Provider Admin Portal.
- 2 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 3 In the left pane, click **External Networks** and click **New**.
- 4 Select a site in which to register the new external network and click **Next**.
- 5 On the **Backing type** page, select **NSX-T Resources (Tier-0 Router)**, select a registered NSX-T Manager to back the network, and click **Next**.

- 6 Enter a name and, optionally, a description for the new external network.
- 7 Select a tier-0 gateway or a VRF gateway to connect to the external network, and click **Next**.
- 8 Configure at least one subnet and click **Next**.
 - a To add a subnet, click **Add**.
 - b Enter the network Classless Inter-Domain Routing (CIDR) settings.
 - c (Optional) Enter the DNS settings.
 - d Configure a static IP pool by adding at least one IP range or IP address.
 - e Click **OK**.
 - f (Optional) To add another subnet, repeat steps 8.a to 8.e.
- 9 Review the network settings and click **Finish**.

What to do next

Use the tier-0 gateway to create an uplink to the external network.

Network Pools

A network pool is a group of undifferentiated networks that is available for use in an organization VDC to create vApp networks and certain types of organization VDC networks.

A network pool is backed by vSphere network resources, such as VLAN IDs or port groups, by NSX Data Center for vSphere resources, or by NSX-T Data Center resources.

VMware Cloud Director uses network pools to create NAT-routed and internal organization VDC networks and all vApp networks. Network traffic on each network in a pool is isolated at layer 2 from all other networks.

Each organization VDC in VMware Cloud Director can have one network pool. Multiple organization VDCs can share a network pool. The network pool for an organization VDC provides the networks created to satisfy the network quota for an organization VDC.

VXLAN Network Pools

Every provider VDC that is backed by NSX Data Center for vSphere includes a VXLAN network pool.

When you create a provider VDC that is backed by NSX Data Center for vSphere, you can associate that provider VDC with an existing VXLAN network pool, or you can create a VXLAN network pool for the provider VDC.

A newly created VXLAN network pool is given a name derived from the name of the containing provider VDC and attached to it at creation. You cannot delete or modify this network pool. If you rename a provider VDC, its VXLAN network pool is automatically renamed.

Note To ensure optimal network performance across your infrastructure, create one VXLAN network pool and associate it with all your provider VDCs upon their creation.

VMware Cloud Director VXLAN networks are based on the IETF VXLAN standard, and provide various benefits.

- Logical networks spanning layer 3 boundaries
- Logical networks spanning multiple racks on a single layer 2
- Broadcast containment
- Higher performance
- Greater scale (up to 16 million network addresses)

For more information about VXLAN networks in a VMware Cloud Director environment, see the *NSX Administration Guide*.

Create a Network Pool Backed by an NSX Data Center for vSphere Transport Zone

To register an NSX Data Center for vSphere transport zone for VMware Cloud Director to use, add a VXLAN-backed network pool.

Prerequisites

Create an NSX Data Center for vSphere transport zone on any vCenter Server registered to VMware Cloud Director. See the *NSX Administration Guide*.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Network Pools** and click **New**.
- 3 Enter a name and, optionally, a description for the new network pool, and click **Next**.
- 4 Select **VXLAN-backed** and click **Next**.
- 5 Select a vCenter Server instance to specify the VXLAN transport zone to be used by this network pool, and click **Next**.
- 6 Select an NSX Data Center for vSphere transport zone to back the new network pool, and click **Next**.

Note To create a universal network pool for cross-virtual data center networking, select a UNIVERSAL_VXLAN type transport zone.

- 7 Review the network pool settings and click **Finish**.

What to do next

Create an organization VDC network that is backed by the network pool or associate the network pool with an organization VDC and create vApp networks.

Geneve Network Pools

Every provider VDC that is backed by NSX-T Data Center includes a Geneve network pool.

Geneve is the network virtualization standard that provides the overlay capability in NSX-T Data Center.

When you create a provider VDC that is backed by NSX-T Data Center, you can associate that provider VDC with an existing Geneve network pool, or you can create a Geneve network pool for the provider VDC.

Note VMware Cloud Director does not support NSX-T Data Center network pools that are backed by VLAN transport zones.

VMware Cloud Director Geneve networks provide a number of benefits.

- Logical networks spanning layer 3 boundaries
- Logical networks spanning multiple racks on a single layer 2
- Broadcast containment
- Higher performance
- Greater scale (up to 16 million network addresses)

Create a Network Pool Backed by an NSX-T Data Center Transport Zone

To register an NSX-T Data Center transport zone for VMware Cloud Director to use, create a Geneve-backed network pool.

Prerequisites

Create an NSX-T Data Center transport zone that is overlay backed.

Note VMware Cloud Director does not support NSX-T Data Center network pools that are backed by VLAN transport zones.

For more information on the transport zone creation and the Generic Network Virtualization Encapsulation, called Geneve overlay, see the *NSX-T Data Center Product Documentation*.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Network Pools** and click **New**.
- 3 Enter a name and, optionally, a description for the new network pool, and click **Next**.
- 4 Select **Geneve-backed** and click **Next**.
- 5 Select an NSX-T Manager instance to provide the transport zone for this network pool, and click **Next**.
- 6 Select an NSX-T transport zone and click **Next**.
- 7 Review the network pool settings and click **Finish**.

What to do next

Create an organization VDC network that is backed by the network pool or associate the network pool with an organization VDC and create vApp networks.

Create a Network Pool Backed by VLAN IDs

To register vSphere VLAN IDs for VMware Cloud Director to use, add a VLAN-backed network pool. A VLAN-backed network pool provides the security, scalability, and performance for organization VDC networks.

Prerequisites

Verify that a range of VLAN IDs and a vSphere distributed switch are available in vSphere. The VLAN IDs must be valid IDs that are configured in the physical switch to which the ESXi servers are connected.

Caution The VLANs must be isolated at the layer 2 level. Failure to isolate properly the VLANs can cause a disruption on the network.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Network Pools** and click **New**.
- 3 Enter a name and, optionally, a description for the new network pool, and click **Next**.
- 4 Select **VLAN-backed** and click **Next**.
- 5 Select a vCenter Server instance to specify the distributed virtual switch to be used by this network pool and click **Next**.
- 6 Enter a VLAN ID range and click **Next**.
- 7 Select a distributed switch for the network pool and click **Next**.
- 8 Review the network pool settings and click **Finish**.

What to do next

Create an organization VDC network that is backed by the network pool or associate the network pool with an organization VDC and create vApp networks.

Create a Network Pool Backed by vSphere Port Groups

To register vSphere port groups for VMware Cloud Director to use, add a network pool backed by port groups. Unlike other types of network pools, a port group-backed network pool does

not require a vSphere distributed switch and can support port groups associated with third-party distributed switches.

Caution The port groups must be isolated from all other port groups at layer 2. The port groups must be physically isolated or must be isolated by using VLAN tags. Failure to isolate properly the port groups can cause a network disruption.

Prerequisites

Verify that one or more port groups are available in your vSphere environment. The port groups must be available on each ESXi host in the cluster, and each port group must use only a single VLAN. Port groups with or without VLAN trunking are supported.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Network Pools** and click **New**.
- 3 Enter a name and, optionally, a description for the new network pool, and click **Next**.
- 4 Select **Portgroup-backed** and click **Next**.
- 5 Select a vCenter Server instance to provide port groups to be used by this network pool, and click **Next**.
- 6 Select one or more port groups and click **Next**.
You can create one network for each port group.
- 7 Review the network pool settings and click **Finish**.

What to do next

Create an organization VDC network that is backed by the network pool or associate the network pool with an organization VDC and create vApp networks.

View the vCenter Server Instances

You can see a list of the vCenter Server instances across all sites in your VMware Cloud Director installation. You can see how VMware Cloud Director uses each vCenter Server instance.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.

Results

A list of all attached vCenter Server instances is displayed. The list contains the following information for each vCenter Server instance.

	Description
Name	The name of the vCenter Server instance in VMware Cloud Director.
Status	The vCenter Server status can be normal, warning, and critical.
State	Activated or deactivated. See Activate or Deactivate a vCenter Server Instance .
Connection	Connected or not to VMware Cloud Director. See Reconnect a vCenter Server Instance .
VC Host	FQDN of the vCenter Server instance.
Version	The vCenter Server version.
Usage	Dedicated vCenter Server instances have enabled tenant access. The provider can use different resource pools of a shared vCenter Server instance across multiple provider VDCs and then allocate those resource pools to different tenants. See Chapter 9 Managing Dedicated vCenter Server Instances .
Cluster Health	Aggregation of the health of all clusters in the vCenter Server instance. When aggregating the health of the cluster, the health of the least healthy cluster is displayed.
Clusters	Number of clusters in the vCenter Server instance.
VMs	Number of VMs in the vCenter Server instance.
Running VMs	Number of running VMs in the vCenter Server instance.
CPU	Amount of actively used virtual CPU as a percentage of the total available vCenter Server CPU.
Memory	Amount of actively used virtual memory as a percentage of the total available vCenter Server memory.
Storage	Amount of actively used virtual storage as a percentage of the total available vCenter Server storage.

Modify vCenter Server Settings

If the connection information for an attached vCenter Server instance changes, or if you want to change its name and description in VMware Cloud Director, or its compute provider scope, you can modify its settings.

You can modify the settings that you configured when adding the vCenter Server instance. See [Add the vCenter Server Instance](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.

2 In the left pane, click **vCenter Server Instances** and click the name of the vCenter Server instance that you want to modify.

3 In the upper right corner of the **vCenter Server Info** section, click **Edit**.

4 (Optional) Edit the instance name and description.

5 (Optional) Edit the compute provider scope for the vCenter Server

The compute provider scope represents compute fault domains, or availability zones which are visible to tenants and where workloads reside. By default, the compute provider scope of a provider virtual data center is inherited from the backing vCenter Server instance. You can differentiate the compute provider scope for the different provider VDCs that are backed by a single vCenter Server instance. For example, you can set the vCenter Server with a compute provider scope **Germany** and you can set the provider VDC with a scope **Munich**.

6 (Optional) Edit the URL for the vCenter Server instance.

7 (Optional) Edit the user name and password for the vCenter Server **administrator** account.

8 (Optional) Turn on or off the **Enabled** toggle.

9 (Optional) Configure the URL of the vCenter Server web client.

10 Click **Save**.

What to do next

If you modified the connection information, you must [Reconnect a vCenter Server Instance](#).

Activate or Deactivate a vCenter Server Instance

Before performing a maintenance or unregistering a vCenter Server instance, you must deactivate the target vCenter Server instance. To provide its resources to virtual data centers in VMware Cloud Director, you must activate the vCenter Server instance.

Procedure

1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.

2 In the left panel, select **vCenter Server Instances**.

3 Click the radio button next to the name of the target vCenter Server instance, and click **Enable** or **Disable**.

4 To confirm, click **OK**.

Reconnect a vCenter Server Instance

If a vCenter Server instance appears as disconnected, or if you modified the connection settings, you can try to reset the connection.

Note During establishing the new connection, the vCenter Server instance is unavailable for operations.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.
- 3 Click the radio button next to the name of the target vCenter Server instance, and click **Reconnect**.
- 4 To confirm, click **OK**.

Refresh a vCenter Server Instance

To update the information in the VMware Cloud Director database about the underlying vCenter Server resources, you must refresh the vCenter Server instance.

Starting with VMware Cloud Director 10.2.2, if you are using Kubernetes, when you refresh a vCenter Server instance, this results in restoring the default firewall policies and NAT rules that block the access to the Tanzu Kubernetes cluster from networks that are outside the organization virtual data center.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.
- 3 Click the radio button next to the name of the target vCenter Server instance, and click **Refresh**.
- 4 To confirm, click **OK**.

Refresh the Storage Policies of a vCenter Server Instance

To update the information in the VMware Cloud Director database about the VM storage policies in the underlying vSphere environment, you must refresh the storage policies of the vCenter Server instance.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.

- 3 Click the radio button next to the name of the target vCenter Server instance, and click **Refresh Policies**.
- 4 To confirm, click **OK**.

Unregister a vCenter Server Instance

To stop using the resources of a vCenter Server instance, you can remove this vCenter Server instance from your VMware Cloud Director installation.

Prerequisites

- Deactivate the vCenter Server instance. See [Activate or Deactivate a vCenter Server Instance](#).
- Delete all provider virtual data centers that use resource pools from this vCenter Server instance. See [Delete a Provider Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.
- 3 Click the radio button next to the name of the target vCenter Server instance, and click **Unregister**.
- 4 To confirm, click **OK**.

Modify NSX Manager Settings

If the connection information for a registered NSX Manager instance changes, or if you want to change its name and description in VMware Cloud Director, you can modify its settings.

You can modify the settings that you configured when adding the NSX Manager instance. See [\(Optional\) Add the Associated NSX Manager Instance](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left pane, click **vCenters** and click the name of the vCenter Server instance that is associated with the target NSX Manager instance.
- 3 In the upper right corner of the **NSX-V Manager Info** section, click **Edit**.
- 4 Modify the NSX Manager hostname and administrator credentials, and click **Save**.

- (Optional) To enable cross-virtual data center networking for the virtual data centers backed by this vCenter Server instance, turn on the toggle, and then enter the control VM properties and a name for the network provider scope.

The control VM properties are used for deploying an appliance on the NSX Manager instance for cross-virtual data center networking components, such as a universal router.

Parameter	Description
Resource Pool Path	The hierarchical path to a specific resource pool in the vCenter Server instance, starting from the cluster, <i>Cluster/Resource_Pool_Parent/Target_Resource</i> . For example, TestbedCluster1/mgmt-rp . As an alternative, you can enter the Managed Object Reference ID of the resource pool. For example, resgroup-1476 .
Datastore Name	The name of the datastore to host the appliance files. For example, shared-disk-1 .
Management Interface	The name of the network in vCenter Server or port group used for the HA DLR management interface. For example, TestbedPG1 .
Network Provider Scope	Corresponds to the network fault domain in the network topologies of the data center groups. For example, boston-fault1 . For information about managing cross-virtual data center groups, see the <i>VMware Cloud Director Tenant Portal Guide</i> .

Modify NSX-T Manager Settings

If the connection information for a registered NSX-T Manager instance changes, or if you want to change its name and description in VMware Cloud Director, you can modify its settings.

You can modify the settings that you configured when adding the vCenter Server instance. See [Register an NSX-T Manager Instance](#).

Procedure

- From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- In the left pane, click **NSX-T Managers** and click the name of the NSX-T Manager instance that you want to modify.
- In the upper right corner of the **General** tab, click **Edit**.
- Edit the NSX-T Manager settings, and click **Save**.

Delete an NSX-T Manager Instance

To stop using the resources of a NSX-T Manager instance, you can remove this vCenter Server instance from your VMware Cloud Director installation.

Prerequisites

Delete all provider virtual data centers that use resources from this NSX-T Manager instance. See [Delete a Provider Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left pane, click **NSX-T Managers**.
- 3 Click the radio button next to the name of the NSX-T Manager instance that you want to remove, and click **Delete**.
- 4 To confirm, click **Delete**.

Configuring and Managing Multisite Deployments

To manage and monitor multiple, geographically distributed VMware Cloud Director installations or server groups and their organizations as single entities, service providers and tenants can use the VMware Cloud Director multisite feature.

Effective Implementation of a Multisite

When you associate two VMware Cloud Director sites, you enable the administration of the sites as a single entity. You also enable organizations at those sites to form associations with each other. When an organization is a member of an association, organization users can use the VMware Cloud Director Tenant Portal to access organization assets at any member site, although each member organization and its assets are local to the site it occupies.

Note To associate sites, you must use the VMware Cloud Director API. The sites must be with the same VMware Cloud Director API version, or one major version apart. For example, you can associate a VMware Cloud Director 10.1 (API version 34.0) site with a VMware Cloud Director site version 10.0, 10.1, 10.2 or 10.2.2, respectively API versions 33.0, 34.0, 35.0, or 35.2.

After you associate two sites, you can use the VMware Cloud Director API or the VMware Cloud Director Tenant Portal to associate organizations that occupy those sites. See the *VMware Cloud Director API Programming Guide* or the [Configure and Manage Multisite Deployments](#) topic in the *VMware Cloud Director Tenant Portal Guide*.

A site or organization can form an unlimited number of associations with a peer, but each association includes exactly two members. Each site or organization must have its own private key. Association members establish a trust relationship by exchanging public keys, which are used to verify signed requests from one member to another.

Each site in an association is defined by the scope of a VMware Cloud Director server group (a group of servers that share a VMware Cloud Director database). Each organization in an association occupies a single site. The organization administrator controls access by organization users and groups to assets at each member site.

Site Objects and Site Associations

The installation or upgrade process creates a `Site` object that represents the local VMware Cloud Director server group. A system administrator whose authority extends to more than one VMware Cloud Director server group can configure those server groups as an association of VMware Cloud Director sites.

Associations of Organizations

After site association is complete, **organization administrators** at any member site can begin associating their organizations.

Note You cannot associate a `System` organization with a tenant organization. The `System` organization at any site can be associated only with the `System` organization at another site.

User and Group Identities

Associations of sites and organizations must agree to use the same identity provider (IDP). User and group identities for all organizations in the association must be managed through this IDP.

Except for the `System` organization, which must use the VMware Cloud Director integrated IDP, associations are free to choose the IDP that works best for them.

Site Access Control for Organization Users and Groups

Organization administrators can configure their IDP to generate user or group access tokens that are valid at all member sites, or valid at only a subset of member sites. While user and group identities must be the same in all member organizations, user and group rights are constrained by the roles those users and groups are assigned in each member organization. Assignment of a role to a user or group is local to a member organization, as are any custom roles you create.

Load Balancer Requirements

Effective implementation of a multisite deployment requires you to configure a load balancer that distributes requests arriving at an institutional endpoint such as `https://vcloud.example.com` to the endpoints for each member of the site association (for example, `https://us.vcloud.example.com` and `https://uk.vcloud.example.com`). If a site has more than one cell, you must also configure a load balancer that distributes incoming requests across all its cells, so that a request to `https://us.vcloud.example.com` can be handled by `https://cell1.us.vcloud.example.com`, `https://cell2.us.vcloud.example.com`, and so on.

Note You must use the global load balancer, in this case `https://vcloud.example.com`, only for UI access. If you develop your own scripts or programs that use the REST API, those calls must target a particular site.

Network Connectivity Requirements

If you want to use the multisite feature, each cell at each site must be able to make REST API requests to the REST API endpoints of all sites. If you use the examples from the Load Balancer Requirements section, `cell1.us.vcloud.example.com` and `cell2.us.vcloud.example.com` must be able to reach the REST API endpoint for `uk.example.com`. The reverse is true for all cells under `uk.example.com`. This means that a cell must also be able to make REST API calls to its own REST API endpoint, so `cell1.us.vcloud.example.com` must be able to make a REST API call to `https://us.vcloud.example.com`.

Making REST API requests to the REST API endpoints of all sites is necessary for REST API fanout. For example, if the UI or an API client makes a multisite request to get a page of organizations from all sites and `cell1.us.vcloud.example.com` handle the request. The cell `cell1` must make a REST API call to get a page of organizations from each site using the REST API endpoint configured for that site. When all sites return their page of organizations, `cell1` collates the results and returns a single page of results containing the data from all other sites.

Sites and Certificates

When a site is associated with other sites, if you update its certificate, you might have to let the other sites know of the change. If you do not let the other sites know about the certificate change, the multisite fanout might be impacted.

If you are replacing a certificate on a site with a valid, well-signed certificate, then you do not need to inform the other sites. Because the certificate is valid and well-signed, the cells at the other sites can continue connecting to it in a secure manner without interruption.

If you are replacing a certificate on a site with a self-signed certificate, or if there is some other problem with the certificate that prevents automatic trust, then other sites need to know. For example, if the certificate expires, you must let the other sites know. At each of the other sites, you must upload the certificate into the **Trusted Certificates** in the Service Provider Admin Portal. See [Import Trusted Certificates](#). When you import the certificate, the site where the certificate is uploaded can trust the site getting the new certificate.

Note You can import these certificates to the Trusted Certificates of the other sites before you install them at the remote site. This ensures no interruptions in communication because both the old certificate and the new certificate are in the Trusted Certificates pool. You do not have to reassociate the sites.

Association Member Status

After you have created an association of sites or organizations, the local system periodically retrieves the status of each remote association member and updates that status in the local site's VMware Cloud Director database. Member status is visible in the `Status` element of an `SiteAssociationMember` OR `OrgAssociationMember`. This element can have one of three values:

ACTIVE

The association has been established by both parties, and communication with the remote party was successful.

ASYMMETRIC

The association has been established at the local site, but the remote site has not yet reciprocated.

UNREACHABLE

An association has been created by both parties, but the remote site is not currently reachable on the network.

The member status "heartbeat" process runs with the identity of the multisite system user, a local VMware Cloud Director user account created in the System organization during VMware Cloud Director installation. Although this account is a member of the System organization, it does not have system administrator rights. It has only a single right, `Multisite: System Operations`, which gives it permission to make a VMware Cloud Director API request that retrieves the status of the remote member of a site association.

Multisite Resource Lists

If you are working with VMware Cloud Director deployments in multiple locations, you can view resource lists that include information about objects from all the connected sites.

To facilitate navigating through vSphere and cloud resources from the Service Provider Admin Portal, starting with version 9.7, VMware Cloud Director introduces multisite resource lists. Starting with version 10.0, VMware Cloud Director supports multisite resource lists that include organizations.

You can access the resource lists through the **vSphere Resources** and the **Cloud Resources** menus.

You can access detailed information about objects from the different sites and also create objects both on the local site and on remote sites.

Multisite vSphere resources lists are supported for vCenter Server instances, NSX-T Manager instances, resource pools, datastores, hosts, distributed switches, port groups, stranded items, and storage policies.

Multisite cloud resources lists are supported for organizations, organization VDCs, organization VDC templates, provider VDCs, cloud cells, edge gateways, external networks, network pools, and VM sizing policies.

Managing Provider Virtual Data Centers

4

After you create a provider virtual data center, you can modify its properties, deactivate or delete it, and manage its storage policies and resource pools.

To create a provider virtual data center, you must use either the Service Provider Admin Portal or the vCloud API. For information about using Service Provider Admin Portal, see [Create a Provider Virtual Data Center](#). For information about using the vCloud API, see the *VMware Cloud Director API Programming Guide*.

This chapter includes the following topics:

- [Activate or Deactivate a Provider Virtual Data Center](#)
- [Delete a Provider Virtual Data Center](#)
- [Edit the General Settings of a Provider Virtual Data Center](#)
- [Merge Provider Virtual Data Centers](#)
- [View the Organization Virtual Data Centers of a Provider Virtual Data Center](#)
- [View the Datastores on a Provider Virtual Data Center](#)
- [View the External Networks on a Provider Virtual Data Center](#)
- [Using Kubernetes with VMware Cloud Director](#)
- [Managing the VM Storage Policies on a Provider Virtual Data Center](#)
- [Managing the Resource Pools on a Provider Virtual Data Center](#)
- [Modify the Metadata for a Provider Virtual Data Center](#)

Activate or Deactivate a Provider Virtual Data Center

To deactivate all existing organization virtual data centers (VDCs) that use the resources of a provider VDC, you can deactivate this provider VDC. You cannot create organization VDCs that use the resources of a deactivated provider VDC.

Running vApps and powered on virtual machines continue to run in the existing organization VDCs backed by this provider VDC, but you cannot create or start additional vApps or virtual machines.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**.
- 3 Click the radio button next to the name of the target provider VDC, and click **Enable** or **Disable**.
- 4 To confirm, click **OK**.

Delete a Provider Virtual Data Center

To remove the resources of a provider virtual data center from VMware Cloud Director, you can delete this provider virtual data center.

The underlying resources in vSphere remain unaffected.

Prerequisites

- Deactivate the target provider virtual data center. See [Activate or Deactivate a Provider Virtual Data Center](#).
- Delete all organization virtual data centers that use resources from this provider virtual data center. See [Delete an Organization Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**.
- 3 Click the radio button next to the name of the provider virtual data center that you want to remove, and click **Delete**.
- 4 To confirm, click **OK**.

Edit the General Settings of a Provider Virtual Data Center

You can change the name and the description of a provider virtual data center. If the backing resource pool supports a higher virtual hardware version, you can upgrade the highest virtual hardware supported by a provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Provider VDCs**, and click the name of the provider virtual data center that you want to modify.
- 3 On the **Configure > General** tab, in the upper right corner, click **Edit**.
- 4 (Optional) Modify the name and the description of the provider virtual data center.

- 5 (Optional) Enter a compute provider scope for the provider virtual data center.

The compute provider scope represents compute fault domains, or availability zones which are visible to tenants and where workloads reside. By default, the compute provider scope of a provider virtual data center is inherited from the backing vCenter Server instance. You can differentiate the compute provider scope for the different provider VDCs that are backed by a single vCenter Server instance. For example, you can set the vCenter Server with a compute provider scope **Germany** and you can set the provider VDC with a scope **Munich**.

- 6 (Optional) From the drop-down menu, select the highest hardware version supported by this provider virtual data center, and click **Save**.

The highest version that you can select depends on the ESXi hosts in the resource pool that backs the provider virtual data center.

Note You can only upgrade the hardware version supported by a provider virtual data center. You cannot downgrade the hardware version. The highest supported virtual machine hardware version in VMware Cloud Director 10.2 is version 17. Hardware version 17 is available when you enable it in the vCenter Server instance on the cluster or data center level.

- 7 Click **Save**.

Merge Provider Virtual Data Centers

To combine the resources of two provider virtual data centers, you can merge these provider virtual data centers into a single provider virtual data center.

Prerequisites

- The target provider virtual data centers belong to the same vCenter Server data center.
- The target provider virtual data centers contain only elastic organization virtual data centers.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**.
- 3 Click the radio button next to the name of the provider virtual data center that you want to expand, and click **Merge**.
- 4 Click the radio button next to the name of the provider virtual data center with which to merge the resources, and click **Merge**.

View the Organization Virtual Data Centers of a Provider Virtual Data Center

You can view a list of the organization virtual data centers that are using resources from a provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Organization VDCs** tab.

Results

The list of the organization virtual data centers that are consuming the resources from this provider virtual data center displays. For each organization VDC, the list includes information about the status, state, allocation model, organization, vCenter Server instance, number of networks, number of vApps, number of storage policies, and number of resource pools.

What to do next

- You can go the organization virtual data center view in the VMware Cloud Director Tenant Portal by clicking the **pop-out** icon () next to the name of the target organization virtual data center.
- By clicking the radio button next to the name of an organization virtual data center, you can perform management operations that are similar to the operations described in [Chapter 6 Managing Organization Virtual Data Centers](#).

View the Datastores on a Provider Virtual Data Center

You can view details about the datastores that provide the storage capacity to a provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Datastores** tab.

A list of all datastores on the provider virtual data center is displayed. The list contains the following information for each datastore.

Title	Description
Name	The name of the datastore
State	Activated or deactivated
Type	The type of file system that the datastore uses, either Virtual Machine File System (VMFS) or Network File System (NFS)

Title	Description
Used	The datastore space occupied by virtual machine files, including log files, snapshots, and virtual disks. When a virtual machine is powered on, the used storage space also includes log files.
Provisioned	The datastore space guaranteed to virtual machines. If any virtual machines are using thin provisioning, some of the provisioned space might not be in use, and other virtual machines can occupy the unused space. This value might be larger than the actual datastore capacity if thin provisioning is used.
Requested Storage	<p>Provisioned storage in use only by VMware Cloud Director objects on the datastore, including:</p> <ul style="list-style-type: none"> ■ Virtual machines provisioned in VMware Cloud Director ■ Catalog items (templates and media) ■ NSX Edges ■ Used and unused memory swap requirements for virtual machines <p>This value does not include storage requested by shadow VMs or intermediate disks in a linked clone tree.</p>
vCenter Server	The vCenter Server instance associated with the datastore.

View the External Networks on a Provider Virtual Data Center

You can view a list of the external networks that are accessible to a provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **External Networks** tab.

Results

You can view a list of the available external networks with information about their gateway CIDR settings and IP pool use.

Using Kubernetes with VMware Cloud Director

By using Kubernetes with VMware Cloud Director, you can provide a multi-tenant Kubernetes service to your tenants.

Container Service Extension

Kubernetes Container Clusters is the Container Service Extension plug-in for VMware Cloud Director. Service providers and tenants must use the Kubernetes Container Clusters plug-in to create Kubernetes clusters. Starting with VMware Cloud Director 10.2, you do not need to download manually the plug-in and upload it to the VMware Cloud Director Service Provider Admin Portal. The plug-in is available in VMware Cloud Director by default, however, you must publish it to tenants to enable them to create Kubernetes clusters.

Both service providers and tenants must use the Container Service Extension version 3.0 to create native and VMware Tanzu Kubernetes Grid Integrated Edition (TKGI) clusters. You must complete the Container Service Extension 3.0 server setup and publish a Container Service Extension native placement policy to one or more organization VDCs.

vSphere with VMware Tanzu in VMware Cloud Director

You can use vSphere with VMware Tanzu in VMware Cloud Director to create provider virtual data centers (VDCs) backed by Supervisor Clusters. A host cluster with enabled vSphere with VMware Tanzu is called a Supervisor Cluster. You can set restrictions on the uses of the resources and limit the available resources, including number of Kubernetes clusters per organization, user, or group. For more information, see [Manage Quotas on the Resource Consumption of an Organization](#).

To use vSphere with VMware Tanzu in VMware Cloud Director, first, you must enable the vSphere with VMware Tanzu functionality on a vSphere 7.0 or later cluster, and configure that cluster as a Supervisor Cluster. See the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation. The vCenter Server instance that you want to use can have both host clusters and Supervisor Clusters.

To create clusters, Tanzu Kubernetes you must publish a provider VDC Kubernetes policy to an organization and apply the organization VDC Kubernetes policy during the creation. Native and TKGI clusters do not use the provider and organization VDC Kubernetes policies.

Kubernetes Cluster Types

- 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. For native Kubernetes cluster deployment, you must set up a Container Service Extension server. See the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- Tanzu Kubernetes clusters - You can use the vSphere with Tanzu 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.

- 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, self-healing, and rolling upgrades for Kubernetes clusters. For more information on TKGI clusters, see the *VMware Tanzu Kubernetes Grid Integrated Edition* documentation.

Workflow for Tanzu Kubernetes Cluster Creation

- 1 Add a vCenter Server 7.0 or later instance with an enabled vSphere with VMware Tanzu functionality to VMware Cloud Director. See [Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance](#).
- 2 Verify the network settings on each Supervisor Cluster to enable them to run Kubernetes workloads.

Important The IP address ranges for the `Ingress CIDRs` and `Services CIDR` parameters must not overlap with IP addresses 10.96.0.0/12 and 192.168.0.0/16 which are the default vSphere values for the `services` and `Pods` parameters. See the configuration parameters for Tanzu Kubernetes clusters information in the *vSphere with Kubernetes Configuration and Management* guide.

Note Starting with VMware Cloud Director 10.2.2, if you modify the network settings of the Supervisor Cluster after the initial setup, you must refresh the vCenter Server instance to adjust the automatic firewall policies and NAT rules that block the access to the Tanzu Kubernetes cluster from outside the organization virtual data center in which the cluster is created.

- 3 Create a provider VDC backed by a Supervisor Cluster. See [Create a Provider Virtual Data Center](#).

Alternatively, you can add a Supervisor Cluster to an existing provider VDC. If you have a vSphere 6.7 or earlier environment, you can also upgrade the environment to version 7.0 and enable vSphere with VMware Tanzu on an existing cluster.

Provider VDCs backed by a Supervisor Cluster appear with a Kubernetes icon next to their name in the grid that lists all provider VDCs.

- 4 (Optional) VMware Cloud Director generates automatically a default provider VDC Kubernetes policy for provider VDCs backed by a Supervisor Cluster. You can create additional provider VDC Kubernetes policies for Tanzu Kubernetes clusters. See [Create a Provider VDC Kubernetes Policy](#).
- 5 [Publish a Provider VDC Kubernetes Policy to an Organization VDC](#) from the **Provider VDCs** tab or [Add an Organization VDC Kubernetes Policy](#) from the **Organization VDCs** tab.

- 6 Publish the Kubernetes Container Clusters plug-in to service providers. See [Publish or Unpublish a Plug-in from an Organization](#). If you want to enable tenants to create Kubernetes clusters, you must publish the Kubernetes Container Clusters plug-in to those organizations. For more information about managing VMware Cloud Director plug-ins, see [Managing Plug-Ins](#).
- 7 If you want to grant tenants the rights to create and manage Tanzu Kubernetes clusters, you must publish the **vmware:tkgcluster Entitlement** rights bundle to any organizations that you want to work with clusters. After sharing the rights bundle, you must add the **Edit: Tanzu Kubernetes Guest Cluster** right to the roles you want to create and modify Tanzu Kubernetes clusters. If you want the users also to delete clusters, you must add the **Full Control: Tanzu Kubernetes Guest Cluster** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).
- 8 Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).
- 9 [Create a Tanzu Kubernetes Cluster](#)

Workflow for Native and TKGI Cluster Creation

- 1 Publish the Kubernetes Container Clusters plug-in to service providers. See [Publish or Unpublish a Plug-in from an Organization](#). If you want to enable tenants to create Kubernetes clusters, you must publish the Kubernetes Container Clusters plug-in to those organizations. For more information about managing VMware Cloud Director plug-ins, see [Managing Plug-Ins](#).
- 2 Set up a Container Service Extension server and publish the Container Service Extension native placement policy or TKGI enablement metadata to the organization VDC. For more information about setting up the CSE server, see the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- 3 If you want to grant tenants the rights to create and manage native clusters, you must publish the **cse:nativeCluster Entitlement** rights bundle to any organizations that you want to work with native clusters. After sharing the rights bundle, you must add the **Edit CSE:NATIVECLUSTER** right to the roles you want to create and modify native clusters. If you want the users also to delete clusters, you must add the **Full Control CSE:NATIVECLUSTER** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).

- 4 If you want to grant tenants the rights to create and manage TKGI clusters, you must publish the **{cse}:PKS DEPLOY RIGHT** to the specific organizations, and add the **{cse}:PKS DEPLOY RIGHT** right to the roles you want to create and manage TKGI clusters. The **{cse}:PKS DEPLOY RIGHT** is created during the Container Service Extension server install.
- 5 For native clusters, grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).
- 6 [Create a Native Kubernetes Cluster](#) or [Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster](#).

Creating a vSphere with VMware Tanzu Cluster

You can use provider VDC and organization VDC Kubernetes policies to create vSphere with VMware Tanzu clusters.

vSphere with VMware Tanzu in VMware Cloud Director

When enabled on a vSphere cluster, vSphere with VMware Tanzu provides the capability to run Kubernetes workloads directly on ESXi hosts and 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.

You can use vSphere with VMware Tanzu in VMware Cloud Director to create provider virtual data centers (VDCs) backed by Supervisor Clusters. A host cluster with enabled vSphere with VMware Tanzu is called a Supervisor Cluster. You can set restrictions on the uses of the resources and limit the available resources, including number of Kubernetes clusters per organization, user, or group. For more information, see [Manage Quotas on the Resource Consumption of an Organization](#).

To use vSphere with VMware Tanzu in VMware Cloud Director, first, you must enable the vSphere with VMware Tanzu functionality on a vSphere 7.0 or later cluster, and configure that cluster as a Supervisor Cluster. See the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation. The vCenter Server instance that you want to use can have both host clusters and Supervisor Clusters.

Tenants can create Tanzu Kubernetes clusters by applying one of the organization VDC Kubernetes policies. System administrators can edit and delete organization VDC Kubernetes policies by using the Service Provider Admin Portal or the VMware Cloud Director Tenant Portal. Native and TKGI clusters do not use the provider and organization VDC Kubernetes policies.

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.

Workflow

- 1 Add a vCenter Server 7.0 or later instance with an enabled vSphere with VMware Tanzu functionality to VMware Cloud Director. See [Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance](#).

- 2 Create a provider VDC backed by a Supervisor Cluster. See [Create a Provider Virtual Data Center](#).

Alternatively, you can add a Supervisor Cluster to an existing provider VDC. If you have a vSphere 6.7 or earlier environment, you can also upgrade the environment to version 7.0 and enable vSphere with VMware Tanzu on an existing cluster.

Provider VDCs backed by a Supervisor Cluster appear with a Kubernetes icon next to their name in the grid that lists all provider VDCs.

- 3 (Optional) VMware Cloud Director generates automatically a default provider VDC Kubernetes policy for provider VDCs backed by a Supervisor Cluster. You can create additional provider VDC Kubernetes policies for Tanzu Kubernetes clusters. See [Create a Provider VDC Kubernetes Policy](#).
- 4 [Publish a Provider VDC Kubernetes Policy to an Organization VDC](#) from the **Provider VDCs** tab or [Add an Organization VDC Kubernetes Policy](#) from the **Organization VDCs** tab.
- 5 Publish the Kubernetes Container Clusters plug-in to service providers. See [Publish or Unpublish a Plug-in from an Organization](#). If you want to enable tenants to create Kubernetes clusters, you must publish the Kubernetes Container Clusters plug-in to those organizations. For more information about managing VMware Cloud Director plug-ins, see [Managing Plug-Ins](#).
- 6 Publish the **vmware:tkgcluster Entitlement** rights bundle to any organizations that you want to work with Tanzu Kubernetes clusters.
- 7 Add the **Edit: Tanzu Kubernetes Guest Cluster** right to the roles that you want to create Tanzu Kubernetes clusters. If you want the users also to delete clusters, you must add the **Full Control: Tanzu Kubernetes Guest Cluster** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).
- 8 Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).
- 9 [Create a Tanzu Kubernetes Cluster](#)

Create a Provider VDC Kubernetes Policy

VMware Cloud Director generates automatically a default provider VDC Kubernetes policy for provider VDCs backed by a Supervisor Cluster. You can create additional provider VDC Kubernetes policies for Tanzu Kubernetes clusters.

Provider VDC and organization VDC Kubernetes policies are necessary only if you want to create or to enable the tenants to create Tanzu Kubernetes clusters. Native and TKGI clusters do not use these Kubernetes policies.

Prerequisites

Verify that you have at least one provider VDC backed by a Supervisor Cluster or add a Supervisor Cluster to an existing provider VDC. See [Using Kubernetes with VMware Cloud Director](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of a provider VDC.
- 3 Under Policies, select **Kubernetes**, and click **New**.

The **Create VDC Kubernetes Policy** wizard appears.

- 4 Enter a name and description for the provider VDC Kubernetes policy and click **Next**.
- 5 Select a resource pool backed by a Kubernetes capable Supervisor Cluster.
- 6 Choose whether you want to reserve CPU and memory for the Kubernetes cluster nodes created in this policy.

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.

- 7 Select CPU and Memory limits for the Kubernetes clusters created under this policy.

When you publish the policy to an organization VDC, the selected limits act as maximums for the newly created organization VDC Kubernetes policy.

- 8 Click **Next**.
- 9 On the **Machine classes** page of the wizard, select one or more VM class types available for this policy, and click **Next**.

The selected machine classes are the only class types available to tenants when you publish the policy to an organization VDC.

- 10 Select one or more storage policies.
- 11 Review your choices and click **Finish**.

What to do next

[Publish a Provider VDC Kubernetes Policy to an Organization VDC](#)

Edit A vSphere Kubernetes Policy

You can edit the settings of provider VDC Kubernetes policies used for the creation of organization VDC Kubernetes policies and Tanzu Kubernetes clusters.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of a provider VDC.
- 3 (Optional) Under Policies, select **Kubernetes**, select the policy you want to publish, and click **Edit**.

The **Edit VDC Kubernetes Policy** wizard appears.

- 4 (Optional) Edit the name and description for the provider VDC Kubernetes policy and click **Next**.
- 5 (Optional) Change the CPU and Memory limits for the Kubernetes clusters created under this policy and click **Next**.

When you publish the policy to an organization VDC, the selected limits act as maximums for the newly created organization VDC Kubernetes policy.

- 6 (Optional) On the **Machine classes** page of the wizard, add one or more VM class types available for this policy, and click **Next**.

The selected machine classes are the only class types available to tenants when you publish the policy to an organization VDC.

- 7 (Optional) Add one or more storage policies.
- 8 Review your choices and click **Save**.

What to do next

[Publish a Provider VDC Kubernetes Policy to an Organization VDC](#)

Publish a Provider VDC Kubernetes Policy to an Organization VDC

To make a provider VDC Kubernetes policy available to tenants, you can publish it to a flex organization VDC. When you publish a provider VDC Kubernetes policy, you create an organization VDC Kubernetes policy that tenants can use to create Kubernetes clusters.

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

You can publish multiple provider VDC Kubernetes policies to a single organization VDC. You can publish a single provider VDC Kubernetes policy multiple times to an organization VDC. 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

- Create a provider VDC backed by a Supervisor Cluster or add a Supervisor Cluster to an existing provider VDC. See [Using Kubernetes with VMware Cloud Director](#).
- Verify that you have at least one flex organization VDC in your environment. See [Create an Organization Virtual Data Center](#).
- 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of a provider VDC.
- 3 Under Policies, select **Kubernetes**, select the policy you want to publish, and click **Publish**.
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 flex organization VDC to which you want to publish the policy and click **Next**.
- 6 Select CPU and Memory limits for the Kubernetes clusters created under this policy.
The maximum limits depend on the CPU and Memory allocations of the organization VDC. When you publish the policy, the selected limits act as maximums for the tenants.

- 7 Choose whether you want to reserve CPU and memory for the 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 publish the policy to an organization VDC.

- 9 Select one or more storage policies.

- 10 Review your choices and click **Publish**.

Results

The information about the published policy appears under the Policies section of the flex organization VDC. 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 Kubernetes clusters. VMware Cloud Director places each 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 Kubernetes clusters in the Supervisor Namespace compete for the resources within these limits.

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 [Using Kubernetes with VMware Cloud Director](#).

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

- Publish the Kubernetes Container Clusters plug-in to any organizations that you want to manage Tanzu Kubernetes 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](#).
- You must publish the **vmware:tkgcluster Entitlement** rights bundle to any organizations that you want to work with clusters. After sharing the rights bundle, you must add the **Edit: Tanzu Kubernetes Guest Cluster** right to the roles you want to create and modify Tanzu Kubernetes clusters. If you want the users also to delete clusters, you must add the **Full Control: Tanzu Kubernetes Guest Cluster** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).
- Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).

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 [Using Kubernetes with VMware Cloud Director](#).

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**.
- To enable the organization VDC for native Kubernetes cluster deployment, set up the Container Service Extension server. See the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.

- Publish the CSE native policy created during the CSE server setup to an organization VDC. To use the UI, see [Add a VM Placement Policy to an Organization VDC](#). Alternatively, you can use the CSE 3.0 CLI to publish policy by running the `vcd cse ovdc enable Organization_VDC_Name --org Organization_Name --native` command.
- You must publish the **cse:nativeCluster Entitlement** rights bundle to any organizations that you want to work with native clusters. After sharing the rights bundle, you must add the **Edit CSE:NATIVECLUSTER** right to the roles you want to create and modify Tanzu Kubernetes clusters. If you want the users also to delete clusters, you must add the **Full Control CSE:NATIVECLUSTER** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).
- Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).

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 [Using Kubernetes with VMware Cloud Director](#).

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

By using the TKGI enablement metadata, you can provide access to the tenants to create TKGI clusters and to access the TKGI-enabled organization VDC. If you want to limit the tenants' ability to create TKGI clusters, you can provide access only to the organization VDC. In this case, the tenants can manage existing TKGI clusters but cannot create new ones.

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**.
- To enable the organization VDC for TKGI Kubernetes cluster deployment, set up the Container Service Extension server. For information about using the CSE CLI to enable an organization VDC for TKGI, see the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- If you want to provide tenant access to TKGI creation and management, you must publish the **{cse}:PKS DEPLOY RIGHT** to the specific organizations, and add the **{cse}:PKS DEPLOY RIGHT** right to the roles you want to create and manage TKGI clusters. The **{cse}:PKS DEPLOY RIGHT** is created during the Container Service Extension server install.

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.

Managing the VM Storage Policies on a Provider Virtual Data Center

You can add, activate, deactivate, and remove VM storage policies from a provider virtual data center (VDC). You can also add, edit, and delete metadata for a VM storage policy on a provider virtual data center.

Starting with VMware Cloud Director 10.2.2, you can limit the allowed entities on a storage policy. See [Edit the Entity Types That a Storage Policy Supports](#) .

Enabling VM Encryption on Storage Policies of a Provider Virtual Data Center

You can add an encryption-enabled storage policy to a provider VDC. You can encrypt VMs and disks by associating a VM or disk with a storage policy that has the VM Encryption capability.

Starting with VMware Cloud Director 10.1, you can improve the security of your data by using VM encryption. 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.

Enabling VM Encryption

To encrypt VMs in VMware Cloud Director, you must configure at least one Key Management Server (KMS) on your vCenter Server instance and associate the VMs and disks with a storage policy that has the VM Encryption capability.

- 1 In vCenter Server, add a KMS cluster. A vCenter Server instance can have multiple KMS clusters. For information about setting up a Key Management Server cluster, see the [Set up the Key Management Server Cluster](#) topic in the *vSphere Security Guide*.
- 2 In vCenter Server, enable encryption on a storage policy. See the [Create an Encryption Storage Policy](#) topic in the *vSphere Security Guide*.

- 3 In the VMware Cloud Director Service Provider Admin Portal , add the encryption-enabled policy to a provider VDC. See [Add a VM Storage Policy to a Provider Virtual Data Center](#).
- 4 In the VMware Cloud Director Service Provider Admin Portal , add the encryption-enabled policy to an organization VDC. See [Add a VM Storage Policy to an Organization Virtual Data Center](#).
- 5 In the VMware Cloud Director Tenant Portal , tenants can associate the VM or disk with a storage policy with enabled VM Encryption.
- 6 To decrypt a VM or disk, tenants can 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 VMs and disks. For more information about roles and rights, see [Predefined Roles and Their Rights](#).

You can view all storage capabilities in the **Capabilities** column under **Resources > vSphere Resources > Storage Policies**. This column displays the VM encryption, tag-based association, vSAN , and IOPS limiting storage capabilities. To view the full list of storage capabilities, expand the row by clicking the arrow on the left side of the storage policy name.

You can also view the storage capability information in the **Storage Policies** tab of a provider VDC.

Add a VM Storage Policy to a Provider Virtual Data Center

You can add a VM storage policy to a provider virtual data center, after which you can configure organization virtual data centers backed by this provider virtual data center to support the added storage policy.

Prerequisites

- Your vSphere administrator created the target VM storage policy. For information about Storage Policy Based Management (SPBM), see the *vSphere Storage* documentation.
- [Refresh the Storage Policies of a vCenter Server Instance.](#)

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Under **Policies**, select **Storage** and click **Add**.
- 4 Select one or more storage policies that you want to add, and click **Add**.

If you select * (**Any**), VMware Cloud Director dynamically adds and removes datastores as they are added to or removed from the datastore clusters of the provider virtual data center.

What to do next

Configure organization virtual data centers backed by the provider virtual data center to support the storage policy. See [Add a VM Storage Policy to an Organization Virtual Data Center](#).

Activate or Deactivate a VM Storage Policy on a Provider Virtual Data Center

After you deactivate a VM storage policy in a provider virtual data center, its organization virtual data centers cannot use this VM storage policy anymore.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the target VM storage policy, and click **Enable** or **Disable**.
- 5 To confirm, click **OK**.

Delete a VM Storage Policy from a Provider Virtual Data Center

You can delete a VM storage policy from a provider virtual data center.

Prerequisites

Deactivate the target VM storage policy. See [Activate or Deactivate a VM Storage Policy on a Provider Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the target VM storage policy, and click **Remove**.
- 5 To confirm, click **Remove**.

Modify the Metadata for a VM Storage Policy on a Provider Virtual Data Center

You can add, edit, and delete metadata for a storage policy on a provider virtual data center.

By using object metadata, you can associate user-defined *name=value* pairs with a storage policy on a provider virtual data center. You can use object metadata in vCloud API query filter expressions.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the target VM storage policy, and click **Metadata**.
- 5 Click **Edit**.
- 6 (Optional) To add a key-value pair, click **Add**, enter a name and a value, and select a type for the new key-value pair.
- 7 (Optional) To edit a key-value pair, enter a new name and a value, and select a new type for the key-value pair.
- 8 (Optional) To remove a key-value pair, in the right end of the row, click the **Delete** icon.
- 9 Click **Save**, and click **OK**.

Enabling the I/O Operations Per Second Setting

You can enable the I/O operations per second (IOPS) setting for a storage policy so that tenants can set per-disk IOPS limits.

Managed read/write performance in physical storage devices and virtual disks is defined in units called IOPS, which measure read/write operations per second. To limit I/O performance, a provider VDC storage policy that includes storage devices with enabled IOPS allocation must back an organization VDC storage policy. Afterwards, a tenant can configure disks that use it to request a specified level of I/O performance. A storage profile configured with IOPS support delivers its default IOPS value to all disks that use it. That includes disks that are not configured to request a specific IOPS value. A hard disk configured to request a specific IOPS value cannot use a storage policy whose maximum IOPS value is lower than the requested value, or a storage policy that is not configured with IOPS support.

Note The actual I/O throughput that the virtual machines see is a combination of block size and IOPS. If the VMs are using different block sizes, their throughput will be different, even if IOPS is limited to the same number. For more information on managing storage I/O resources, see the *vSphere Resource Management* guide.

VMware Cloud Director IOPS Storage Policy

With this option, there are default IOPS settings that you can edit. You can set limits on IOPS per disk or IOPS per storage policy. You can set IOPS limits per disk based on the disk size in GB so that you grant larger disks more IOPS. Tenants can set custom IOPS on a disk within these limits. You can use IOPS limiting with or without IOPS capacity considerations for placement.

You cannot enable IOPS on a storage policy backed by a Storage DRS cluster.

- 1 If you want VMware Cloud Director to consider IOPS when placing disks on datastores, in vCenter Server, add IOPS capacities to all datastores associated with the storage policy you want to modify.
- 2 If you want VMware Cloud Director to consider IOPS when placing disks on datastores, in vCenter Server, create a storage policy that uses the datastores with added IOPS capacities.
- 3 By using the VMware Cloud Director Service Provider Admin Portal or the VMware Cloud Director API, add the storage policy to one or more provider VDCs.
- 4 By using the Service Provider Admin Portal or the VMware Cloud Director API, publish the storage policy to one or more organization VDCs. The organization VDCs to which you publish the storage policy inherit the policy's IOPS settings.
- 5 If you want to edit the inherited storage policy IOPS settings, use the Service Provider Admin Portal or VMware Cloud Director API to update the organization VDC storage policy.

This policy type appears as a `VCD/IOPS` capability of the storage policy.

vCenter Server IOPS Storage Policy

This option has one IOPS setting for all disks using this policy. You cannot edit this setting in VMware Cloud Director. Tenants cannot set custom IOPS on disks using these policies. This option does not provide IOPS scaling depending on the sizes of the disks or load balancing across datastores.

- 1 In vCenter Server, create a VC-IOPS enabled storage policy with custom reservation, limit, and shares.
- 2 In vCenter Server or the VMware Cloud Director Service Provider Admin Portal, assign the disk to the storage policy.

This policy type appears as a `vSphere/IOPS` capability of the storage policy. When the source or target VM has the `vSphere/IOPS` capability, you cannot create fast-provisioned VMs.

Setting IOPS on a Disk in vCenter Server

To change the IOPS setting, in vCenter Server, manually update the IOPS on the disk. You cannot edit these IOPS settings in VMware Cloud Director.

Enabling IOPS Limiting on an Existing Storage Policy

Note You cannot enable VMware Cloud Director IOPS limiting on a policy that already has the `vSphere/IOPS` capability on it.

- Enable IOPS limiting on a `VCD/IOPS` storage policy:
 - a If you want VMware Cloud Director to consider IOPS capacities when placing disks on datastores, in vCenter Server, add IOPS capacities to all datastores associated with the storage policy you want to modify.
 - b If you want VMware Cloud Director to consider IOPS capacities when placing disks on datastores, by using the VMware Cloud Director Service Provider Admin Portal or the VMware Cloud Director API, ensure that the corresponding provider VDC storage policy reports the IOPS capacity as non-zero.
 - c By using the VMware Cloud Director Service Provider Admin Portal or VMware Cloud Director API, update the organization VDC storage policy to enable the `VCD/IOPS` capability and to set the maximum IOPS value, default IOPS value, and so on.
- Enable IOPS limiting on a `vSphere/IOPS` storage policy in vCenter Server.

When you enable IOPS limiting for an organization VDC storage policy, tenants can use the VMware Cloud Director Tenant Portal to set per-disk IOPS limits.

Edit the Provider VDC Storage Policy Settings

You can change the I/O operations per second (IOPS) settings of a provider VDC storage policy. By default, the organization VDCs to which the policy is published inherit the provider VDC storage policy settings.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Under **Policies**, select **Storage**.
- 4 Click the radio button next to the target storage policy, and click **Edit Settings**.
- 5 If you want to limit the I/O operations per second, turn on the **IOPS Limiting Enabled** toggle.
- 6 If you want IOPS to be considered during placement, turn on the **Impact Placement** toggle.
 If the **Impact Placement** toggle is turned on, VMware Cloud Director provides IOPS load balancing across datastores. When you set IOPS settings for a disk, VMware Cloud Director considers datastores with enough IOPS capacity for the selected disk. If the **Impact Placement** toggle is turned off, you do not need to set IOPS capacities per datastore and you can use Storage DRS clusters.
- 7 Configure the maximum and default IOPS settings and click **Save**.

Results

The new storage policy settings apply to all organization VDCs to which this policy is published.

Edit the Entity Types That a Storage Policy Supports

Starting with VMware Cloud Director 10.2.2, if you do not want a provider VDC storage policy to support certain types of VMware Cloud Director entities, you can edit and limit the list of entities associated with the policy.

When you create a provider VDC storage policy, by default, it supports all available entity types. The default entity types are:

- Virtual machines
- Named disks
- Catalog media
- vApp and VM templates
- Tanzu Kubernetes clusters
- Edge gateways

You can limit the entity types that a storage policy supports to one or more types from this list. When you create an entity, only the storage policies that support its type are available. For example, if you want to create a catalog, the only storage policies that appear are the ones that support catalog media, vApp templates, or both. If an entity uses a storage policy and you remove the type of entity from the list of supported entity types, the entity continues to use the storage policy but you cannot make any changes to it without selecting a new storage policy.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the target storage policy, and click **Edit Supported Types**.
- 5 From the **Supports Entity Types** drop-down menu, select **Select Specific Entities**.
- 6 Select the entities that you want the storage policy to support, and click **Save**.

What to do next

- [Add a VM Storage Policy to an Organization Virtual Data Center](#)
- Users with the **Supported Storage Entity Type: Manage** right can use the VMware Cloud Director OpenAPI to add or remove entity types from the list of available types for all storage policies. For example, you can add or remove Runtime Defined Entities (RDEs) to the list. For more information about creating extensions that provide additional VMware Cloud Director capabilities to the tenants, see [Chapter 14 Managing Defined Entities](#).

VMware Cloud Director automatically applies the changes to the storage policies that support all entities. You cannot remove entities that are selected specifically in one or more storage policies.

Managing the Resource Pools on a Provider Virtual Data Center

You can add, activate, deactivate, and detach secondary resource pools from a provider virtual data center. You cannot deactivate or detach the primary resource pool on a provider virtual data center.

Add a Resource Pool to a Provider Virtual Data Center

You can add one or more secondary resource pools to a provider virtual data center, so that its Pay-As-You-Go and Allocation Pool organization virtual data centers can expand.

When compute resources are backed by multiple resource pools, they can expand to accommodate more virtual machines.

You can add resource pools backed by vSphere clusters that are optimally configured for hosting NSX edges that have VLAN uplinks. VMware Cloud Director can use metadata to indicate that the system must place organization VDC Edge Gateways in resource pools backed by those clusters. For more information, see VMware Knowledge Base Article <https://kb.vmware.com/kb/2151398>.

Prerequisites

Your vSphere administrator created the target secondary resource pool in the vCenter Server instance that backs the primary resource pool of the provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 On the **Resource Pools** tab, click **Add**.
- 4 Select the resource pool you want to add, and click **Add**.
If you want to use vSphere with VMware Tanzu, select a Supervisor Cluster. VMware Cloud Director displays a Kubernetes icon next to resource pools backed by a Supervisor Cluster.
- 5 If you select a resource pool or cluster that is backed by a Supervisor Cluster, to establish a trust relationship with the Kubernetes control plane, you must trust the Kubernetes control plane certificate.
- 6 If you want to add an additional resource pool, repeat [Step 1](#) to [Step 5](#).

Results

VMware Cloud Director adds the resource pool for the provider virtual data center to use, making elastic all Pay-As-You-Go and Allocation Pool organization virtual data centers backed by the provider virtual data center.

VMware Cloud Director also adds a System VDC resource pool beneath the new resource pool. This resource pool is used for the creation of system resources such as NSX edge VMs and VMs that serve as a template for linked clones.

Important Do not edit or delete the System VDC resource pool.

Activate or Deactivate a Resource Pool on a Provider Virtual Data Center

When you deactivate a resource pool, the memory and compute resources of the resource pool are no longer available to the provider virtual data center.

Processes that are already in progress do not stop using resources from the deactivated resource pool.

Note You cannot deactivate the primary resource pool on a provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.

- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Resource Pools** tab.
- 4 Click the radio button next to the target resource pool, and click **Enable** or **Disable**.
- 5 To confirm, click **OK**.

Detach a Resource Pool from a Provider Virtual Data Center

If a provider virtual data center has more than one resource pool, you can detach a secondary resource pool from the provider virtual data center. You cannot detach the primary resource pool from the provider virtual data center.

Prerequisites

- Deactivate the target resource pool on the provider virtual data center. See [Activate or Deactivate a Resource Pool on a Provider Virtual Data Center](#).
- Redeploy any networks that are affected by the deactivated resource pool.
- Redeploy any edge gateways that are affected by the deactivated resource pool.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Resource Pools** tab.
- 4 Click the radio button next to the target resource pool, and click **Detach**.
- 5 To confirm, click **OK**.

Modify the Metadata for a Provider Virtual Data Center

You can add, edit, and delete metadata for a provider virtual data center.

By using object metadata, you can associate user-defined *name=value* pairs with a provider virtual data center. You can use object metadata in vCloud API query filter expressions.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 On the **Configure > Metada** tab, in the upper right corner, click **Edit**.
- 4 (Optional) To add a key-value pair, click **Add**, enter a name and a value, and select a type for the new key-value pair.

- 5 (Optional) To edit a key-value pair, enter a new name and a value, and select a new type for the key-value pair.
- 6 (Optional) To remove a key-value pair, in the right end of the row, click the **Delete** icon.
- 7 Click **Save**, and click **OK**.

Managing Organizations

5

The VMware Cloud Director Service Provider Admin Portal allows you to create, configure, and manage VMware Cloud Director organizations.

Use VMware Cloud Director Service Provider Admin Portal to manage organizations, set policies to determine how users consume resources allocated to an organization, and manage publishing and sharing of catalogs.

This chapter includes the following topics:

- [Understanding Leases](#)
- [Create an Organization](#)
- [Activate or Deactivate an Organization](#)
- [Delete an Organization](#)
- [Configure Catalogs for an Organization](#)
- [Configure Policies for an Organization](#)
- [Migrate Tenant Storage](#)
- [Manage Quotas on the Resource Consumption of an Organization](#)

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 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.

Create an Organization

You can create a new organization from the VMware Cloud Director Service Provider Admin Portal.

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.

- a From the left panel, select **Organizations**.

The list of existing organizations displays in a grid view.

- 2 Click **New**.

The **New Organization** dialog opens.

- 3 Enter the following values.

Option	Description
Organization name	The unique identifier that forms the URL for accessing the Tenant Portal of the organization.
Organization full name	Full name of the organization.
Description	An optional description for the organization.

- 4 Click the **Create** button to complete the creation.

Activate or Deactivate an Organization

Deactivating an organization prevents users from logging in to the organization and terminates the sessions of users that are currently logged in. Running vApps in the organization continue to run.

As a **system administrator**, you can allocate resources, add networks, and so on, even after you deactivate an organization.

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.

- a From the left panel, select **Organizations**.

The list of existing organizations displays in a grid view.

- 2 Click the radio button next to the name of the organization and click **Enable** or **Disable**.

Delete an Organization

Delete an organization to permanently remove it from VMware Cloud Director.

Prerequisites

Before you can delete an organization, you must deactivate it and delete all organization virtual data centers, templates, media files, and vApps in the organization.

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.
 - a From the left panel, select **Organizations**.
The list of existing organizations displays in a grid view.
- 2 Click the radio button next to the name of the organization and click **Delete**.
- 3 To confirm, click **Yes**.

Configure Catalogs for an Organization

You can configure how an organization shares its service catalogs.

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.
 - a From the left panel, select **Organizations**.
The list of existing organizations displays in a grid view.
- 2 Select an organization, and under the **Configure** tab, select **Catalog**.
- 3 To change the sharing and publishing settings, click **Edit**.

Option	Description
Sharing	Allows organization administrators to share this organization's catalogs with other organizations in this instance of VMware Cloud Director. If you do not select this option, organization administrators are still able to share catalogs within the organization.
Allow publishing to external catalogs	Allows organization administrators to publish catalogs to organizations outside of this instance of VMware Cloud Director.
Allow subscribing to external catalogs	Allows organization administrators to subscribe to catalogs outside of this instance of VMware Cloud Director.

Configure Policies for an Organization

Leases, quotas, and limits constrain the ability of organization users to consume storage and processing resources. You can modify these settings to prevent users from depleting or monopolizing an organization's resources.

Prerequisites

See [Understanding Leases](#).

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.

- a From the left panel, select **Organizations**.

The list of existing organizations displays in a grid view.

- 2 Select an organization and select the **Policies** tab.
- 3 To edit the leases, quotas, resource limits, and password policies for the organization, click **Edit**.
- 4 Configure vApp leases with the following settings.

Option	Description
Maximum runtime lease	How long vApps can run before they are automatically stopped.
Runtime expiry action	How expired running vApps are processed. Suspending a vApp, suspends all its virtual machines and preserves their current state by writing the memory to disk. Power off immediately stops all its virtual machines and child vApps.
Maximum storage lease	How long stopped vApps are available before being automatically cleaned up.
Storage cleanup	How vApps are processed after being stopped and cleaned up.

- 5 Configure vApp template leases with the following settings.

Option	Description
Maximum storage lease	How long vApp templates are available before being automatically cleaned up.
Storage cleanup	How expired vApp templates are processed after being cleaned up.

- 6 Configure quotas with the following settings.

Option	Description
All VMs quota	Total number of available VMs a user can store in this organization.
Running VMs quota	Total number of VMs a user can power on in this organization.

7 Configure limits with the following settings.

Option	Description
Number of resource intensive operations per user	Type the maximum number of simultaneous resource intensive operations per user, or select Inherit System Limit .
Number of resource intensive operations to be queued per user	Type the maximum number of queued resource intensive operations per user, or select Inherit System Limit .
Number of resource intensive operations per organization	Type the maximum number of simultaneous resource intensive operations per organization, or select Inherit System Limit .
Number of resource intensive operations to be queued per organization	Type the maximum number of queued resource intensive operations per organization, or select Inherit System Limit .
Number of simultaneous connections per VM	Type the maximum number of simultaneous console connections per virtual machine, or select Inherit System Limit .
Number of virtual data centers per organization	Type the maximum number of organization virtual data centers per organization, or select Inherit System Quota .

8 Configure password policies with the following settings.

Option	Description
Account lockout enabled	Enable user account lockout after a number of invalid login attempts.
Invalid logins before lockout	Number of invalid login attempts before the user account is locked.
Account lockout interval	Period during which a locked user account cannot log in.

Migrate Tenant Storage

You can migrate all vApps, independent disks, and catalog items of one or more organizations from one or more datastores to different datastores.

Before you decommission a datastore, you must migrate all the items stored on that datastore to a new datastore. You might also want to migrate an organization to a new datastore that has more storage capacity or uses a newer storage technology such as VMware vSAN.

Important Tenant storage migration is a resource-intensive operation that can run for a long time, especially when there are many assets to migrate. For more information about migrating tenant storage, see <https://kb.vmware.com/kb/2151086>.

Prerequisites

- Determine the storage policies used by the organization VDCs of the target organizations. See [Add a VM Storage Policy to an Organization Virtual Data Center](#).
- For each storage policy containing a source datastore that you want to migrate, verify that there is at least one destination datastore to which to migrate. You can create destination datastores or use existing ones. For further information about determining the datastores in the storage policies used by the target organizations, see the *vSphere Storage* documentation.

Procedure

- 1 Log in to the VMware Cloud Director Service Provider Admin Portal as a **system administrator** or with a role that has the **Organization: Migrate Tenant Storage** right.
- 2 Start the **Migrate Tenant Storage** wizard.
 - Under **Cloud Resources**, select **Organizations** and click **Migrate Tenant Storage**.
 - Under **vSphere Resources**, select **Datastores** and click **Migrate Tenant Storage**.
- 3 Select one or more organizations with storage items that you want to migrate, and click **Next**.
- 4 Select one or more source datastores to migrate, and click **Next**.
The wizard lists all datastores in the system.
- 5 Select one or more destination datastores, and click **Next**.
- 6 Review the **Ready to Complete** page, and click **Finish** to begin the migration.

Manage Quotas on the Resource Consumption of an Organization

You can manage the overall resource consumption limit of an organization. You can add, edit, and remove the organization's quotas on VMs, Tanzu Kubernetes clusters, CPU, memory, or storage.

For information about limiting the resources available to users or groups, see [Manage the Resource Quotas of a User](#) or [Manage the Resource Quotas of a Group](#) .

Prerequisites

[Create an Organization](#)

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 From the left panel, select **Organizations**.
- 3 Select the name of the organization for which you want to put on a quota.
- 4 Under the **Configure** section, select **Quotas** .

Organizations do not have any quotas by default.

- 5 Click **Edit**.
- 6 Modify the quota for the selected organization.

You can add, edit, or remove quotas on the number of Tanzu Kubernetes clusters, all or running VMs in the organization, consumed CPU, memory, and storage. Select **Unlimited** if you want the organization to have unlimited resources of the selected type.

- 7 Click **Save**.

Managing Organization Virtual Data Centers

6

To provide resources to an organization, you create one or more organization virtual data centers (VDCs) for this organization. After you create an organization VDC, you can modify its properties, deactivate or delete it, and manage its allocation model, storage, and network settings.

This chapter includes the following topics:

- [Understanding Allocation Models](#)
- [Understanding VM Sizing and VM Placement Policies](#)
- [Using Kubernetes with VMware Cloud Director](#)
- [Create an Organization Virtual Data Center](#)
- [Activate or Deactivate an Organization Virtual Data Center](#)
- [Delete an Organization Virtual Data Center](#)
- [Managing Virtual Data Center Templates](#)
- [Modify the Name and the Description of an Organization Virtual Data Center](#)
- [Modify the Allocation Model Settings of an Organization Virtual Data Center](#)
- [Modifying the Storage Settings of an Organization Virtual Data Center](#)
- [Edit the Network Settings of an Organization Virtual Data Center](#)
- [Configuring Cross-Virtual Data Center Networking](#)
- [Modify the Metadata for an Organization Virtual Data Center](#)
- [View the Resource Pools of an Organization Virtual Data Center](#)
- [Managing the Distributed Firewall on an Organization Virtual Data Center](#)

Understanding Allocation Models

An allocation model determines how and when the allocated provider virtual data center (VDC) compute and memory resources are committed to the organization VDC.

The following table shows the vSphere resource distribution settings at the virtual machine (VM) or resource pool level based on the organization VDC allocation model.

	Flex Allocation Model	Elastic Allocation Pool Model	Non-Elastic Allocation Pool Model	Pay-As-You-Go Model	Reservation Pool Model
Elastic	Based on the organization VDC configuration.	Yes	No	Yes	No
vCPU Speed	If a VM CPU limit is not defined in a VM sizing policy, vCPU speed might impact the VM CPU limit within the VDC.	Impacts the number of running vCPUs in the Organization VDC.	Not Applicable	Impacts VM CPU Limit	Not Applicable
Resource Pool CPU Limit	Organization VDC CPU limit apportioned based on the number of VMs in the resource pool.	Organization VDC CPU allocation	Organization VDC CPU allocation	Unlimited	Organization VDC CPU allocation
Resource Pool CPU Reservation	Organization VDC CPU reservation is apportioned based on the number of vCPUs in the resource pool. Organization VDC CPU reservation equals the organization VDC CPU allocation times the CPU guarantee.	Sum of powered on VMs and equals the CPU guarantee times the vCPU speed, times the number of vCPUs.	Organization VDC CPU allocation times the CPU guarantee	None, expandable	Organization VDC CPU allocation
Resource Pool Memory Limit	Organization VDC memory limit is apportioned based on the number of VMs in the resource pool.	Unlimited	Organization VDC RAM allocation	Unlimited	Organization VDC RAM allocation
Resource Pool Memory Reservation	Organization VDC RAM reservation is apportioned based on the number of VMs in the resource pool. The organization VDC RAM reservation equals the organization VDC RAM allocation times the RAM guarantee.	Sum of RAM guarantee times vRAM of all powered-on VMs in the resource pool. The resource pool RAM reservation is expandable.	Organization VDC RAM allocation times the RAM guarantee	None, expandable	Organization VDC RAM allocation
VM CPU Limit	Based on the VM sizing policy of the VM.	Unlimited	Unlimited	vCPU speed times the number of vCPUs	Custom
VM CPU Reservation	Based on the VM sizing policy of the VM.	0	0	Equals the CPU speed times the vCPU speed, times the number of vCPUs.	Custom

	Flex Allocation Model	Elastic Allocation Pool Model	Non-Elastic Allocation Pool Model	Pay-As-You-Go Model	Reservation Pool Model
VM RAM Limit	Based on the VM sizing policy of the VM.	Unlimited	Unlimited	vRAM	Custom
VM RAM Reservation	Based on the VM sizing policy of the VM.	0	Equals vRAM times RAM guarantee plus RAM overhead.	Equals vRAM times RAM guarantee plus RAM overhead.	Custom

Converting a Legacy VDC Allocation Model to a Flex Allocation Model

You add a VM placement and VM sizing policy to a VDC with an elastic allocation pool model, non-elastic allocation pool model, pay-as-you-go model, or reservation pool model. If the VM placement or VM sizing policy is not compatible with the existing VDC allocation model, you can decide to convert the VDC to a flex organization VDC.

VM Policy Compliance

Legacy VDC conversion does not cause VM non-compliance. If an administrator changes the VM compute values or VM group membership of a VM directly in the vCenter Server instance, a VM can become non-compliant with the assigned VM sizing or VM placement policy. A VM can also become non-compliant if a user with the necessary privileges changes the VM reservation and limit values by using the vCloud API. If there is a non-compliant VM, VMware Cloud Director Tenant Portal UI displays a warning message. The tenant can see detailed information about the cause for the non-compliance and can make the VM compliant again, which reapplies the policies to the VM.

Suggested Use of the Allocation Models

Each allocation model can be used for different levels of performance control and management. The following table contains information about the suggested use of each allocation model.

Allocation model	Suggested use
Flex allocation model	With the flex allocation model, you can achieve a fine-grained performance control at the workload level. By using the flex allocation model, VMware Cloud Director system administrators can manage the elasticity of individual organization VDCs. The flex allocation model uses policy-based management of workloads. With the flex allocation model, cloud providers can have a better control over memory overhead in an organization VDC and can enforce a strict burst capacity use for tenants.
Allocation pool allocation model	Use the allocation pool allocation model for long lived, stable workloads, where tenants subscribe to a fixed compute resource consumption and where cloud providers can predict and manage the compute resource capacity. The allocation pool allocation model is optimal for workloads with diverse performance requirements. With the allocation pool allocation model, all workloads share the allocated resources from the resource pools of vCenter Server. Regardless if you activate or deactivate elasticity, tenants receive a limited amount of compute resources. With the allocation pool allocation model, cloud providers activate or deactivate the elasticity at the system level and the setting applies to all allocation pool organization VDCs. If you use the non-elastic allocation pool allocation, the organization VDC pre-reserves the VDC resource pool and tenants can overcommit vCPUs but cannot overcommit any memory. If you use the elastic pool allocation, the organization VDC does not pre-reserve any compute resources and capacity can span through multiple clusters. Cloud providers manage the overcommitment of physical compute resources and tenants cannot overcommit vCPUs and memory.
Pay-as-you-go	Use the pay-as-you-go model when you do not have to allocate compute resources in vCenter Server upfront. Reservation, limit, and shares are applied on every workload that tenants deploy in the VDC. With the pay-as-you-go allocation model, every workload in the organization VDC receives the same percentage of the configured compute resources reserved. To VMware Cloud Director, the CPU speed of every vCPU for every workload is the same and you can only define the CPU speed at the organization VDC level. From a performance perspective, because you cannot change reservation settings of individual workloads, every workload receives the same preference. Pay-as-you-go allocation model is optimal for tenants that need workloads with different performance requirements to run within the same organization VDC. Because of the elasticity, the pay-as-you-go model is suitable for generic, short lived workloads that are part of autoscaling applications. With pay-as-you-go, tenants can match spikes in compute resources demand within an organization VDC.
Reservation pool	Use the reservation pool allocation model when you need a fine-grained control over the performance of workloads that are running in the organization VDC. From a cloud provider perspective, the reservation pool allocation model requires an upfront allocation of all compute resources in vCenter Server. The reservation pool allocation model is not elastic. The reservation pool allocation model is optimal for workloads that run on hardware that is dedicated to a specific tenant. In such cases, tenant users can manage use and overcommitment of compute resources.

Flex Allocation Model

Starting with VMware Cloud Director 9.7, **system administrators** can create organization virtual data centers (VDC) by using the flex allocation model. With the combination of flex allocation and VM sizing policies, **system administrators** can control CPU and RAM consumption at both the VDC and the individual virtual machine (VM) levels. The flex allocation model supports all allocation configurations that are available in the existing allocation models.

In VMware Cloud Director 10.0 and later, all non-flex organization VDCs can be converted into flex VDCs.

When creating a flex organization VDC, **system administrators** control the following parameters of the organization VDC:

Parameter	Description
Elasticity	Activate or deactivate the elastic pool feature.
Include VM Memory Overhead	Include or exclude memory overhead in this VDC. When set to true, you might not be able to use the full capacity of the VDC because the memory overhead of every powered-on VM is also taken from the available capacity of the VDC. When set to false, the memory overhead is taken from the provider VDC and not from the allocated capacity of the VDC.
CPU allocation	The amount of CPU allocated to this VDC in MHz or GHz. The CPU allocation defines the CPU capacity of the VDC. The total CPU used by all VMs running in the VDC cannot exceed this value.
CPU limit	The CPU limit defines the CPU quota of a VDC. In most cases, the CPU limit is equal to the allocated CPU capacity of the VDC. Sometimes, you might be required not to allocate any CPU to the VDC, as in the pay-as-you-go model. In this case, you must set a quota on the overall CPU consumption by setting the CPU allocation to zero and the CPU limit to a non-zero value. You might also use this setting to allow an unlimited CPU quota. If set to unlimited, the backing resource pools of the VDC in vCenter Server get unlimited CPU.
CPU resources guaranteed	The percentage of CPU allocation that is physically reserved for the VDC.
vCPU speed	The default vCPU speed for VMs in the VDC.
Memory allocation	The amount of memory allocated to this VDC in MB or GB. This parameter defines the total memory capacity of the VDC. The total configured memory by all VMs running in the VDC cannot exceed this value.
Memory resources guaranteed	The percentage of memory allocation that is physically reserved for the VDC.
Maximum number of VMs	The maximum number of VMs in the VDC.

As a **VMware Cloud Director system administrator**, you can configure a flex organization VDC to be elastic or non-elastic. When flex organization VDCs have the elastic pool feature enabled, the organization VDC spans and uses all resource pools that are associated with its provider VDC. In VMware Cloud Director 9.7, if you convert a non-elastic organization VDC to an elastic organization VDC, you cannot convert the same organization VDC back to a non-elastic.

The flex allocation model supports the capabilities of VM sizing policies without any constraints that other allocation models have. In the flex allocation model, the VM compute resource allocation depends on the VM sizing policies. If you do not define a VM sizing policy for an organization VDC, the compute resource allocation depends on the organization VDC allocation model. Using the combination of the flex allocation model and the organization VM sizing policies, a single organization VDC can accommodate VMs that use configuration that is common for all other allocation models. For more information, see [Understanding VM Sizing and VM Placement Policies](#).

To create a flex organization VDC, you can use the VMware Cloud Director Service Provider Admin Portal or vCloud API. For information about vCloud API, see *VMware Cloud Director API Programming Guide*.

Allocation Pool Allocation Model

With the allocation pool allocation model, a percentage of the resources you allocate from the provider virtual data center (VDC) are committed to the organization VDC. You can specify the percentage for both CPU and memory. This percentage is known as the percentage guarantee factor, and it allows you to overcommit resources.

As a system administrator, you can configure allocation-pool organization VDCs to be elastic or non-elastic. Elasticity is a global setting that affects all allocation-pool organization VDCs. See [Modify General System Settings](#).

By default, allocation-pool organization VDCs have an elastic allocation pool enabled. Systems upgraded from VMware Cloud Director 5.1 that have allocation-pool organization VDCs with virtual machines spanning multiple resource pools have elastic allocation pool enabled by default.

When allocation-pool VDCs have the elastic allocation pool feature enabled, the organization VDC spans and uses all resource pools associated with its provider VDC. As a result, vCPU frequency is now a mandatory parameter for an allocation pool.

Set the vCPU frequency and percentage guarantee factor in such a way that enough virtual machines can be deployed on the organization VDC without CPU being a bottleneck factor.

When a virtual machine is created, the placement engine places it on a provider VDC resource pool that best fits the requirements of the virtual machine. A subresource pool is created for this organization VDC under the provider VDC resource pool, and the virtual machine is placed under that subresource pool.

When the virtual machine powers on, the placement engine checks the provider VDC resource pool to ensure that it still can power on the virtual machine. If not, the placement engine moves the virtual machine to a provider VDC resource pool with sufficient resources to run the virtual machine. A subresource pool for the organization VDC is created if one does not exist.

The subresource pool is configured with sufficient resources to run the new virtual machine. The subresource pool's memory reservation is increased by the virtual machine's configured memory size times the percentage guarantee factor for the organization VDC. The subresource pool's CPU reservation is increased by the number of vCPU configured for the virtual machine times the vCPU specified at the organization VDC level times the percentage guarantee factor for CPU set at the organization VDC level. If the elastic allocation pool feature is enabled, the subresource pool's memory limit is increased by the virtual machine's configured memory size, and the subresource pool's CPU limit is increased by the number of vCPUs that the virtual machine is configured with times the vCPU frequency specified at the organization VDC level. The virtual machine is reconfigured to set its memory and CPU reservation to zero and the virtual machine placement engine places the virtual machine on a provider VDC resource pool.

With the elastic allocation pool allocation model, the limits are monitored and managed by VMware Cloud Director only. If the elastic feature is deactivated, the resource pool limit is set additionally.

The benefits of the allocation-pool model are that a virtual machine can take advantage of the resources of an idle virtual machine on the same subresource pool. This model can take advantage of new resources added to the provider VDC.

In rare cases, a virtual machine is switched from the resource pool it was assigned at creation to a different resource pool at power-on because of a lack of resources on the original resource pool. This change might involve a minor cost to move the virtual machine disk files to a new resource pool.

When the elastic allocation pool feature is deactivated, the behavior of allocation-pool organization VDCs is similar to the allocation pool model in VMware Cloud Director 1.5. In this model, the vCPU frequency is not configurable. Overcommitment is controlled by setting the percentage of resources guaranteed.

By default, in an allocation pool VDC, virtual machines obtain their reservation, limit, and shares settings from the settings of the VDC. To create or reconfigure a virtual machine with custom resource allocation settings for both CPU and memory, you can use the vCloud API. See *VMware Cloud Director API Programming Guide*.

Pay-As-You-Go Allocation Model

With the pay-as-you-go allocation model, resources are committed only when users create vApps in the organization VDC. You can specify a percentage of resources to guarantee, which allows you to overcommit resources. You can make a pay-as-you-go organization VDC elastic by adding multiple resource pools to its provider VDC.

Resources committed to the organization are applied at the virtual machine level.

When a virtual machine is powered on, if the original resource pool cannot accommodate the virtual machine, the placement engine checks the resource pool and assigns the virtual machine to another resource pool. If a subresource pool is not available for the resource pool, VMware Cloud Director creates one with an infinite limit and zero rate. The virtual machine's rate is set to its limit times its committed resources, and the virtual machine placement engine places the virtual machine on a provider VDC resource pool.

The benefit of the pay-as-you-go model is that it can take advantage of new resources added to the provider VDC.

In rare cases, a virtual machine is switched from the resource pool it was assigned at creation to a different resource pool at power-on because of a lack of resources on the original resource pool. This change might involve a minor cost to move the virtual machine disk files to a new resource pool.

In the pay-as-you-go model, no resources are reserved ahead of time, so a virtual machine might fail to power on if there are not enough resources. Virtual machines operating under this model cannot take advantage of the resources of idle virtual machines on the same subresource pool, because resources are set at the virtual machine level.

By default, in a pay-as-you-go VDC, virtual machines obtain their reservation, limit, and shares settings from the settings of the VDC. To create or reconfigure a virtual machine with custom resource allocation settings for both CPU and memory, you can use the vCloud API. See *VMware Cloud Director API Programming Guide*.

Reservation Pool Allocation Model

With the reservation pool allocation model, all the resources you allocate are immediately committed to the organization VDC. Users in the organization can control the overcommitment by specifying reservation, limit, and priority settings for individual virtual machines.

Because only one resource pool and one subresource pool are available in this model, the placement engine does not reassign a virtual machine's resource pool when it is powered on. The virtual machine's rate and limit are not modified.

With the reservation pool model, sources are always available when needed. This model also offers fine control over the virtual machine rate, limit, and shares, which can lead to optimal use of the reserved resources if you plan carefully. For information about configuring virtual machine resource allocation settings in reservation pool VDCs, see the *vCloud Air - Virtual Private Cloud OnDemand User's Guide*.

In this model, reservation is always done at the primary cluster. If sufficient resources are not available to create an organization VDC on the primary cluster, the organization VDC creation fails.

Other limitations of this model are that it is not elastic and organization users might set nonoptimal shares, rates, and limits on virtual machines, leading to underuse of resources.

Understanding VM Sizing and VM Placement Policies

You can control the virtual machine (VM) resource allocation and placement on a specific cluster or host by using VM sizing policies and VM placement policies.

VMware Cloud Director **system administrators** create and manage VM sizing policies at a global level and can publish individual policies to one or more organization VDCs. For VMware Cloud Director 10.2.1 and earlier, you can create and manage VM placement policies for each provider VDC separately, because a VM placement policy is scoped at the provider VDC level. Starting with VMware Cloud Director 10.2.2, you can include more than one provider VDC in the scope of a VM placement policy. In addition, starting with version 10.2.2, if a user saves a vApp as a vApp template to a catalog, the template includes also the placement and sizing policies of the original vApp as unmodifiable tagged policies.

When you publish a policy to an organization VDC, the policy becomes available to the users in the organization. When creating and managing virtual machines in the organization VDC, tenants can assign the available policies to the virtual machines. Tenants and users in the organization VDC cannot look into the specific configuration of a VM placement policy or a VM sizing policy.

VM placement and sizing policies are a mechanism for cloud providers to define and offer differentiated levels of service, for example, a CPU intensive profile or a high memory usage profile. If you publish multiple VM placement and VM sizing policies to an organization VDC, tenant users can select between all custom policies and the default policy when creating and managing VMs in the organization VDC. The System Default policy is auto-generated for every VDC. You can delete the System Default policy in the VDC and mark another custom policy as the default. The default policy does not define any values and allows all virtual machine configurations.

VM placement policy

A VM placement policy defines the placement of a virtual machine on a host or group of hosts. It is a mechanism for **cloud provider administrators** to create a named group of hosts within a provider VDC. The named group of hosts is a subset of hosts within the provider VDC clusters that might be selected based on any criteria such as performance tiers or licensing. Starting with VMware Cloud Director 10.2.2, you can expand the scope of a VM placement policy to more than one provider VDC.

A VM placement policy defines VM-host affinity rules that directly impact the placement of tenant workloads. Administrators define or expose named host groups by using VM groups in vCenter Server. A VM group has a direct affinity to a host group and represents the host group to which it has the affinity.

You define the VM placement policy at the provider VDC level. A VM placement policy includes the following attributes:

- Name (must be unique in the provider VDC)
- Description
- A set of one or more VM groups selected from the underlying clusters in the provider VDC. You can select one VM group per cluster

A VM placement policy is optional during the creation of a virtual machine and a tenant can assign only one VM placement policy to a virtual machine.

When a tenant creates a virtual machine in the organization VDC and selects the VM placement policy, VMware Cloud Director adds the virtual machine to the VM group or VM groups that are referenced in the policy. As a result, VMware Cloud Director creates the virtual machine on the appropriate host.

A VM placement policy can have zero or one VM group from each cluster. For example, the VM placement policy *oracle_license* can comprise VM groups *oracle_license1* and *oracle_license2*, where VM group *oracle_license1* belongs to cluster *oracle_cluster1*, and VM group *oracle_license2* belongs to cluster *oracle_cluster2*.

When you assign a VM placement policy to a virtual machine, the placement engine adds this virtual machine to the corresponding VM group of the cluster on which it resides. For example, if you select to deploy a virtual machine on cluster *oracle_cluster1* and assign the VM placement policy *oracle_license* to this virtual machine, the placement engine adds the virtual machine to VM group *oracle_license1*.

VM sizing policy

A VM sizing policy defines the compute resource allocation for virtual machines within an organization VDC. The compute resource allocation includes CPU and memory allocation, reservations, limits, and shares.

With VM sizing policies, VMware Cloud Director **system administrators** can control the following aspects of compute resources consumption at the virtual machine level:

- Number of vCPUs and vCPU clock speed
- Amount of memory allocated to the virtual machine
- Memory and CPU reservation, limit, and shares
- Extra Configurations.

The `extraConfigs` API parameter represents a mapping between a key and value pairs that are applied as extra configuration values on a virtual machine. You can create a policy with extra configurations only by using the vCloud API. Existing extra configurations appear in the Service Provider Admin Portal UI under **Extra Configurations** in the detailed VM sizing policy view.

You define the VM sizing policies at a global level. For more information about the VM sizing policy attributes, see [Attributes of VM Sizing Policies](#).

VMware Cloud Director generates a default VM sizing policy for all VDCs. The default VM sizing policy contains only a name and description, and all remaining policy attributes are empty.

You can also define another VM sizing policy as the default policy for an organization VDC. The default VM sizing policy controls the resource allocation and consumption of the virtual machines that tenants create in the organization VDC, unless a tenant assigns another specific VM sizing policy to the virtual machine.

To limit the maximum compute resources that tenants can allocate to individual virtual machines within an organization VDC, cloud providers can define a maximum VM sizing policy. When assigned to an organization VDC, the maximum VM sizing policy acts as an upper bound for the compute resource configuration for all virtual machines within the organization VDC. The maximum VM sizing policy is not available to tenant users when creating a virtual machine. When you define a VM sizing policy as the maximum policy, VMware Cloud Director copies internally the content of the policy and uses the copied content as the maximum VM sizing policy. As a result, the organization VDC does not depend on the initially used VM sizing policy.

By using VM sizing policies, cloud providers can restrict the compute resources consumption for all virtual machines within an organization VDC to, for example, three predefined sizes, *Small Size*, *Medium Size*, and *Large Size*. The workflow is the following.

- 1 A **system administrator** creates three VM sizing policies with the following attributes.

Name	Attributes
Small Size	<ul style="list-style-type: none"> ■ Description: Small-sized VM policy ■ Name: Small Size ■ Memory: 1024 ■ Number of vCPUs: 1
Medium Size	<ul style="list-style-type: none"> ■ Description: Medium-sized VM policy ■ Name: Medium Size ■ Memory: 2048 ■ Number of vCPUs: 2
Large Size	<ul style="list-style-type: none"> ■ Description: Large-sized VM policy ■ Name: Large Size ■ Memory: 4096 ■ Number of vCPUs: 4

- 2 Publish the new VM sizing policies to an organization VDC.
- 3 Optionally define one of the VM sizing policies as a default VM sizing policy for the organization VDC.

The available policy operations for cloud providers are the following:

- To define the placement of a virtual machine on a host or group of hosts, create a placement policy. See [Create a VM Placement Policy within a Provider VDC](#).
- To control the physical compute resource allocation for tenant workloads, create a sizing policy. See [Create a VM Sizing Policy](#).
- Publish a VM placement or VM sizing policy to one or more organization VDCs. See [Add a VM Placement Policy to an Organization VDC](#)
- Set a VM placement or VM sizing policy as default.
- Edit a VM placement policy and a VM sizing policy. You can only edit the name and description of the policy in the VMware Cloud Director UI.
- Unpublish a VM placement or VM sizing policy from an organization VDC.
- Delete a VM placement or VM sizing policy. See [Delete a VM Placement Policy](#) and [Delete a VM Sizing Policy](#).

Users that have the **ORG_VDC_MANAGE_COMPUTE_POLICIES** right can create, update, and publish VM placement and VM sizing policies.

The following table lists the available VM sizing policy and VM placement policy operations for tenant users.

Table 6-1. VM Sizing Policy and VM Placement Policy Operations for Tenant Users

Operation	Description
Assign a policy to a virtual machine during a virtual machine creation.	<p>Tenant users that are authorized to create virtual machines in an organization VDC can optionally assign VM sizing and VM placement policies to virtual machines by using the VMware Cloud Director Tenant Portal. As a result, the parameters defined in the VM sizing policy control the CPU and memory consumption of the virtual machine. Assigning a VM placement or sizing policy is not a requirement for tenants during a virtual machine creation. If a tenant does not explicitly select a VM sizing policy to assign to a virtual machine, the default VM sizing is applied to the virtual machine. If you do not create any VM placement policy, the VM placement policy option is not visible to the tenants. If the tenant selects a placement policy that has sizing information, the VM sizing policy option becomes hidden to the tenant. You can create a VM placement policy with sizing information only by using the vCloud API. If there is only one VM sizing policy, the VM sizing policy option is not visible to the tenants.</p> <p>When the system administrator sets the vCPU Count, Cores Per Socket, and Memory attributes in a VM sizing policy, if a tenant selects the policy, these values are shown, but not editable.</p>
Assign a policy to an existing virtual machine.	<p>Tenant users that are authorized to manage virtual machines in an organization VDC can assign or change the VM sizing and VM placement policies of an existing virtual machine using the VMware Cloud Director Tenant Portal. When a tenant changes the VM placement policy, the virtual machine moves to a new host as per the VM-host affinity rule defined in the new VM placement policy. When a tenant changes a VM sizing policy, the system reconfigures the virtual machine to consume compute resources as specified in the new VM sizing policy.</p>

The workflow for working with VM placement and VM sizing policies is the following.

- 1 A **system administrator** creates one or more VM placement policies. See [Create a VM Placement Policy within a Provider VDC](#).
- 2 A **system administrator** creates one or more VM sizing policies. See [Create a VM Sizing Policy](#).

The name of a VM sizing policy is unique in a single VMware Cloud Director site. The name of a VM placement policy is unique within the provider VDC scope of the policy.

- 3 A **system administrator** publishes the VM placement and VM sizing policies to one or more organization VDCs. See [Add a VM Placement Policy to an Organization VDC](#).

Publishing a VM placement policy makes it available for tenant users in the organization VDCs during virtual machine creation and virtual machine editing.

- 4 When creating or updating a virtual machine, tenants can use the vCloud API or the VMware Cloud Director Tenant Portal to assign a VM sizing policy and a VM placement policy to a virtual machine.

Create a VM Placement Policy within a Provider VDC

A VM placement policy is a VDC compute policy that contains a reference to a provider VDC policy. Starting with VMware Cloud Director 10.2.2, you can add multiple provider VDCs to the

scope of a VM placement policy. You can use a VM placement policy to define the placement of a VM on a specific host, group of hosts, or a cluster.

Starting with VMware Cloud Director 10.2.2 a VM placement policy can contain a reference to one or more provider VDC policies. When you create a placement policy from within a provider VDC, the policy references only the selected provider VDC. You can include more provider VDCs in the scope of a VM placement policy by editing it or you can create a placement policy from the **VM Placement Policies** tab to include more than one provider VDC in its scope. See [Edit a VM Placement Policy](#) and [Create a Global VM Placement Policy](#).

Prerequisites

- Verify that you have at least one provider VDC in your environment.
- Verify that you have at least one VM group in your environment.

A VM group is a collection of VMs that you can link to a host group with positive affinities. Through a positive affinity rule, you cause the placement of a group of VMs on a specific host. You can create a VM group through the vCenter Server UI or the VMware Cloud Director API.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**.
- 3 Click a provider VDC from the list.
- 4 Click the **VM Placement Policies** tab and click **New**.
- 5 (Optional) On the **What is VM Placement policy** page of the wizard, select the check box to stop showing the VM placement policy information.
- 6 Click **Next**.
- 7 Enter a name for the VM placement policy and, optionally, a description.
- 8 Select the VM groups or logical VM groups to which you want the VM to be linked and click **Next**.

When you select more than one logical group, if a tenant applies this policy to a VM, the VM becomes a member of all the VM groups included in the selected logical VM groups. The VM is conditioned to a combination of all the affinities that apply to the VMs in these groups. Starting with VMware Cloud Director 10.2.2, you can select simultaneously VM groups and logical groups.

You can create an inline logical VM group by selecting one VM group per cluster. This logical VM group does not have a name and can be used only for the selected VM Placement policy.

- 9 Review the VM placement policy settings and click **Finish**.

What to do next

- [Create a VM Sizing Policy](#).

- [Add a VM Placement Policy to an Organization VDC.](#)
- Starting with VMware Cloud Director 10.2.2, you can [Edit a VM Placement Policy.](#)
- [Delete a VM Placement Policy.](#)

Create a Global VM Placement Policy

Starting with VMware Cloud Director 10.2.2 a VM placement policy can contain a reference to one or more provider VDC policies. You can use a VM placement policy to define the placement of a VM on a specific host, group of hosts, or one or more clusters.

When you create a placement policy from within a provider VDC, the policy references only the selected provider VDC. See [Create a VM Placement Policy within a Provider VDC.](#) Starting with VMware Cloud Director 10.2.2, you can include more provider VDCs in the scope of a VM placement policy by editing it, or you can create a global placement policy.

Prerequisites

- Verify that you have at least one provider VDC in your environment.
- Verify that you have at least one VM group in your environment.

A VM group is a collection of VMs that you can link to a host group with positive affinities. Through a positive affinity rule, you cause the placement of a group of VMs on a specific host. You can create a VM group through the vCenter Server UI or the VMware Cloud Director API.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 From the left panel, select **VM Placement Policies**, and click **New**.
- 3 (Optional) On the **What is VM Placement policy** page of the wizard, select the check box to stop showing the VM placement policy information.
- 4 Click **Next**.
- 5 Enter a name for the VM placement policy and, optionally, a description.
- 6 Select the VM groups and logical VM groups to which you want the VM to be linked and click **Next**.

You can select one VM group per cluster.

When you select more than one logical group, if a tenant applies this policy to a VM, the VM becomes a member of all the VM groups included in the selected logical VM groups. The VM is conditioned to a combination of all the affinities that apply to the VMs in these groups. Starting with VMware Cloud Director 10.2.2, you can select simultaneously VM groups and logical groups.

You can create an inline logical VM group by selecting one VM group per cluster. This logical VM group does not have a name and can be used only for the selected VM Placement policy.

7 Review the VM placement policy settings and click **Finish**.

What to do next

- [Create a VM Sizing Policy](#).
- [Add a VM Placement Policy to an Organization VDC](#).
- Starting with VMware Cloud Director 10.2.2, you can [Edit a VM Placement Policy](#).
- [Delete a VM Placement Policy](#).

Edit a VM Placement Policy

Starting with VMware Cloud Director 10.2.2, you can edit and change the scope of a VM placement policy.

Prerequisites

[Create a Global VM Placement Policy](#)

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 From the left panel, select **VM Placement Policies**.
- 3 Select a VM placement policy, and click **Edit**.
- 4 (Optional) On the **What is VM Placement policy** page of the wizard, select the check box to stop showing the VM placement policy information.
- 5 Click **Next**.
- 6 Edit the name for the VM placement policy and, optionally, the description.
- 7 Edit the VM groups and logical VM groups to which you want the VM to be linked and click **Next**.

You can select one VM group per cluster. You cannot deselect clusters that are currently in use, for example, when you publish the placement policy to an organization VDC.

- 8 Review the VM placement policy settings and click **Finish**.

What to do next

- [Create a VM Sizing Policy](#).
- [Add a VM Placement Policy to an Organization VDC](#).
- [Delete a VM Placement Policy](#).

Add a VM Placement Policy to an Organization VDC

When you create a VM placement policy, it is not visible to tenants. You can publish a VM placement policy to an organization VDC to make it available to tenants.

Publishing a VM placement policy to an organization VDC makes the policy visible to tenants. For VMware Cloud Director 10.2.2 and later, to publish a placement policy to an organization VDC, you must first include its backing provider VDC in the scope of the VM placement policy by [Create a Global VM Placement Policy](#) or [Edit a VM Placement Policy](#). The tenant can select the policy when they create a new standalone VM or a VM from a template, edit a VM, add a VM to a vApp, and create a vApp from a vApp template. You cannot delete a VM placement policy that is available to tenants.

Prerequisites

- Verify that you have at least one organization VDC in your environment. See [Create an Organization Virtual Data Center](#).
- Verify that you have at least one VM placement policy. See [Create a VM Placement Policy within a Provider VDC](#). For VMware Cloud Director 10.2.2 and later, you can create a global placement policy that contains a reference to one or more provider VDC policies. See [Create a Global VM Placement Policy](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Select an organization VDC and click the **VM Placement Policies** tab.
- 4 Click **Add**.
- 5 Select the VM placement policies that you want to add to the organization VDC and click **OK**.

What to do next

- Select a policy and click **Remove** to unpublish the policy.
- Select a VM placement policy and click **Set as default** to make that policy appear as the default choice for the tenants during a VM and vApp creation and VM edit. If there is more than one VM placement policy published for an organization VDC, the tenant can select a different policy from the default one.

Delete a VM Placement Policy

If a VM placement policy is not published to tenants, you can delete it from the provider VDC.

Prerequisites

- Verify that you have at least one VM placement policy in your environment.
- Verify that the VM placement policy is not added to an organization VDC. You cannot delete VM placement policies that are available to tenants.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.

- 2 In the left panel, select **Provider VDCs**.
- 3 Click a provider VDC from the list.
- 4 Click the **VM Placement Policies** tab and select a VM placement policy.
- 5 Click **Delete**.

Attributes of VM Sizing Policies

When you create a virtual machine (VM) sizing policy, you can specify a subset of all available attributes. The only mandatory attribute is the VM sizing policy name.

There are two types of parameters in a VM sizing policy.

- Individual VM sizing configuration - You preconfigure the specified RAM, vCPU count, and cores per socket for the VMs under the current policy.
- Constraints on the maximum resources - You preconfigure a limitation for consumption of memory and CPU by a single VM under the current policy.

The following table lists all attributes that you can define within a VM sizing policy.

Table 6-2. VDC Compute Policy Attributes

VDC Compute Policy Attribute	API Parameter	Description
Name	name	Mandatory parameter that is used as an identifier for the VM sizing policy.
Description	description	Represents a short description of the VM sizing policy.
vCPU Speed	cpuSpeed	Defines the vCPU speed of a core in MHz or GHz.
vCPU Count	cpuCount	Defines the number of vCPUs configured for a VM. This is a VM hardware configuration. When a tenant assigns the VM sizing policy to a VM, this count becomes the configured number of vCPUs for the VM.
Cores Per Socket	coresPerSocket	The number of cores per socket for a VM. This is a VM hardware configuration. The number of vCPUs that is defined in the VM sizing policy must be divisible by the number of cores per socket. If the number of vCPUs is not divisible by the number of cores per socket, the number of cores per socket becomes invalid.
CPU Reservation Guarantee	cpuReservationGuarantee	Defines how much of the CPU resources of a VM are reserved. The allocated CPU for a VM equals the number of vCPUs times the vCPU speed in MHz. The value of the attribute ranges between 0 and one. Value of 0 CPU reservation guarantee defines no CPU reservation. Value of 1 defines 100% of CPU reserved.
CPU Limit	cpuLimit	Defines the CPU limit in MHz or GHz for a VM. If not defined in the VDC compute policy, CPU limit is equal to the vCPU speed multiplied by the number of vCPUs.

Table 6-2. VDC Compute Policy Attributes (continued)

VDC Compute Policy Attribute	API Parameter	Description
CPU Shares	cpuShares	<p>Defines the number of CPU shares for a VM.</p> <p>Shares specify the relative importance of a VM within a virtual data center. If a VM has twice as many shares of CPU as another VM, it is entitled to consume twice as much CPU when these two virtual machines are competing for resources.</p> <p>If not defined in the VDC compute policy, normal shares are applied to the VM.</p>
Memory	memory	<p>Defines the memory configured for a VM in MB or GB. This is a VM hardware configuration.</p> <p>When a tenant assigns the VM sizing policy to a VM, the VM receives the amount of memory defined by this attribute.</p>
Memory Reservation Guarantee	memoryReservationGuarantee	<p>Defines the reserved amount of memory that is configured for a VM.</p> <p>The value of the attribute ranges between 0 and 100%.</p>
Memory Limit	memoryLimit	<p>Defines the memory limit in MB or GB for a VM.</p> <p>If not defined in the VM sizing policy, memory limit is equal to the allocated memory for the VM.</p>
Memory Shares	memoryShares	<p>Defines the number of memory shares for a VM.</p> <p>Shares specify the relative importance of a VM within a virtual data center. If a VM has twice as many shares of memory as another VM, it is entitled to consume twice as much memory when these two virtual machines are competing for resources.</p> <p>If not defined in the VDC compute policy, normal shares are applied to the VM.</p>
Extra Configuration s	extraConfigs	<p>Represents a mapping between a key and value pairs that are applied as extra configuration values on a VM.</p> <p>You can create a policy with extra configurations only through the vCloud API. Existing extra configurations appear in the Service Provider Admin Portal UI under Extra Configurations in the detailed VM sizing policy view.</p>

Create a VM Sizing Policy

You can create a VM sizing policy to make available to tenants predefined CPU and memory consumption constraints that they can apply to individual VMs in an organization VDC.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **VM Sizing Policies**.
- 3 Click **New**.
- 4 Enter a name for the VM sizing policy, and optionally a description.
- 5 Click **Next**.
- 6 On the **CPU** page, select the CPU allocation settings that you want to apply to the policy and click **Next**.

- 7 Select the memory allocation settings that you want to apply to the policy and click **Next**.
- 8 Review the VM sizing policy settings and click **Finish**.

What to do next

- After you create a VM sizing policy, you can edit only the VM sizing policy name and description. See [Edit a VM Sizing Policy](#).
- [Add a VM Sizing Policy to an Organization VDC](#).
- [Create a VM Placement Policy within a Provider VDC](#).

Add a VM Sizing Policy to an Organization VDC

When you create a VM sizing policy, it is not visible to tenants. You can publish VM sizing policy to an organization VDC to make it available to tenants.

Publishing a VM sizing policy to an organization VDC makes the policy visible to tenants. The tenant can select the policy when they create a new standalone VM or a VM from a template, edit a VM, add a VM to a vApp, and create a vApp from a vApp template. You cannot delete a VM sizing policy that is available to tenants.

Prerequisites

- Verify that you have at least one organization VDC in your environment. See [Create an Organization Virtual Data Center](#).
- Verify that you have at least one VM sizing policy. See [Create a VM Sizing Policy](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Select an organization VDC and click the **VM Sizing Policies** tab.
- 4 Click **Add**.
- 5 Select the VM sizing policies that you want to add to the organization VDC and click **OK**.

What to do next

- Select a policy and click **Remove** to unpublish the policy.
- Select a VM sizing policy and click **Set as default** to make that policy appear as the default choice for the tenants during a VM and vApp creation and VM edit. If there is more than one VM sizing policy published for an organization VDC, the tenant can select a different policy from the default one.

Edit a VM Sizing Policy

After you create a VM sizing policy, you can edit only its name and description. Editing the CPU and memory parameters is not supported.

Prerequisites

- Verify that you have at least one organization VDC in your environment. See [Create an Organization Virtual Data Center](#).
- Verify that you have at least one VM sizing policy. See [Create a VM Sizing Policy](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **VM Sizing Policies**.
- 3 Click the name of the VM sizing policy you want to edit.
- 4 To edit the name and description of the policy, click **Edit**.
- 5 Click **Save**.

What to do next

[Add a VM Sizing Policy to an Organization VDC](#)

Delete a VM Sizing Policy

You can delete VM sizing policies that are not published to tenants.

Prerequisites

- Verify that you have at least one VM sizing policy in your environment.
- Verify that the VM sizing policy is not added to an organization VDC. You cannot delete VM sizing policies that are available to tenants.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **VM Sizing Policies**.
- 3 Select a VM sizing policy and click **Delete**.

Using Kubernetes with VMware Cloud Director

By using Kubernetes with VMware Cloud Director, you can provide a multi-tenant Kubernetes service to your tenants.

Container Service Extension

Kubernetes Container Clusters is the Container Service Extension plug-in for VMware Cloud Director. Service providers and tenants must use the Kubernetes Container Clusters plug-in to create Kubernetes clusters. Starting with VMware Cloud Director 10.2, you do not need to download manually the plug-in and upload it to the VMware Cloud Director Service Provider Admin Portal. The plug-in is available in VMware Cloud Director by default, however, you must publish it to tenants to enable them to create Kubernetes clusters.

Both service providers and tenants must use the Container Service Extension version 3.0 to create native and VMware Tanzu Kubernetes Grid Integrated Edition (TKGI) clusters. You must complete the Container Service Extension 3.0 server setup and publish a Container Service Extension native placement policy to one or more organization VDCs.

vSphere with VMware Tanzu in VMware Cloud Director

You can use vSphere with VMware Tanzu in VMware Cloud Director to create provider virtual data centers (VDCs) backed by Supervisor Clusters. A host cluster with enabled vSphere with VMware Tanzu is called a Supervisor Cluster. You can set restrictions on the uses of the resources and limit the available resources, including number of Kubernetes clusters per organization, user, or group. For more information, see [Manage Quotas on the Resource Consumption of an Organization](#).

To use vSphere with VMware Tanzu in VMware Cloud Director, first, you must enable the vSphere with VMware Tanzu functionality on a vSphere 7.0 or later cluster, and configure that cluster as a Supervisor Cluster. See the *vSphere with Kubernetes Configuration and Management* guide in the vSphere documentation. The vCenter Server instance that you want to use can have both host clusters and Supervisor Clusters.

To create clusters, Tanzu Kubernetes you must publish a provider VDC Kubernetes policy to an organization and apply the organization VDC Kubernetes policy during the creation. Native and TKGI clusters do not use the provider and organization VDC Kubernetes policies.

Kubernetes Cluster Types

- 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. For native Kubernetes cluster deployment, you must set up a Container Service Extension server. See the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- Tanzu Kubernetes clusters - You can use the vSphere with Tanzu 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.

- 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, self-healing, and rolling upgrades for Kubernetes clusters. For more information on TKGI clusters, see the *VMware Tanzu Kubernetes Grid Integrated Edition* documentation.

Workflow for Tanzu Kubernetes Cluster Creation

- 1 Add a vCenter Server 7.0 or later instance with an enabled vSphere with VMware Tanzu functionality to VMware Cloud Director. See [Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance](#).
- 2 Verify the network settings on each Supervisor Cluster to enable them to run Kubernetes workloads.

Important The IP address ranges for the `Ingress CIDRs` and `Services CIDR` parameters must not overlap with IP addresses 10.96.0.0/12 and 192.168.0.0/16 which are the default vSphere values for the `services` and `Pods` parameters. See the configuration parameters for Tanzu Kubernetes clusters information in the *vSphere with Kubernetes Configuration and Management* guide.

Note Starting with VMware Cloud Director 10.2.2, if you modify the network settings of the Supervisor Cluster after the initial setup, you must refresh the vCenter Server instance to adjust the automatic firewall policies and NAT rules that block the access to the Tanzu Kubernetes cluster from outside the organization virtual data center in which the cluster is created.

- 3 Create a provider VDC backed by a Supervisor Cluster. See [Create a Provider Virtual Data Center](#).

Alternatively, you can add a Supervisor Cluster to an existing provider VDC. If you have a vSphere 6.7 or earlier environment, you can also upgrade the environment to version 7.0 and enable vSphere with VMware Tanzu on an existing cluster.

Provider VDCs backed by a Supervisor Cluster appear with a Kubernetes icon next to their name in the grid that lists all provider VDCs.

- 4 (Optional) VMware Cloud Director generates automatically a default provider VDC Kubernetes policy for provider VDCs backed by a Supervisor Cluster. You can create additional provider VDC Kubernetes policies for Tanzu Kubernetes clusters. See [Create a Provider VDC Kubernetes Policy](#).
- 5 [Publish a Provider VDC Kubernetes Policy to an Organization VDC](#) from the **Provider VDCs** tab or [Add an Organization VDC Kubernetes Policy](#) from the **Organization VDCs** tab.

- 6 Publish the Kubernetes Container Clusters plug-in to service providers. See [Publish or Unpublish a Plug-in from an Organization](#). If you want to enable tenants to create Kubernetes clusters, you must publish the Kubernetes Container Clusters plug-in to those organizations. For more information about managing VMware Cloud Director plug-ins, see [Managing Plug-Ins](#).
- 7 If you want to grant tenants the rights to create and manage Tanzu Kubernetes clusters, you must publish the **vmware:tkgcluster Entitlement** rights bundle to any organizations that you want to work with clusters. After sharing the rights bundle, you must add the **Edit: Tanzu Kubernetes Guest Cluster** right to the roles you want to create and modify Tanzu Kubernetes clusters. If you want the users also to delete clusters, you must add the **Full Control: Tanzu Kubernetes Guest Cluster** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).
- 8 Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).
- 9 [Create a Tanzu Kubernetes Cluster](#)

Workflow for Native and TKGI Cluster Creation

- 1 Publish the Kubernetes Container Clusters plug-in to service providers. See [Publish or Unpublish a Plug-in from an Organization](#). If you want to enable tenants to create Kubernetes clusters, you must publish the Kubernetes Container Clusters plug-in to those organizations. For more information about managing VMware Cloud Director plug-ins, see [Managing Plug-Ins](#).
- 2 Set up a Container Service Extension server and publish the Container Service Extension native placement policy or TKGI enablement metadata to the organization VDC. For more information about setting up the CSE server, see the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- 3 If you want to grant tenants the rights to create and manage native clusters, you must publish the **cse:nativeCluster Entitlement** rights bundle to any organizations that you want to work with native clusters. After sharing the rights bundle, you must add the **Edit CSE:NATIVECLUSTER** right to the roles you want to create and modify native clusters. If you want the users also to delete clusters, you must add the **Full Control CSE:NATIVECLUSTER** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).

- 4 If you want to grant tenants the rights to create and manage TKGI clusters, you must publish the **{cse}:PKS DEPLOY RIGHT** to the specific organizations, and add the **{cse}:PKS DEPLOY RIGHT** right to the roles you want to create and manage TKGI clusters. The **{cse}:PKS DEPLOY RIGHT** is created during the Container Service Extension server install.
- 5 For native clusters, grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).
- 6 [Create a Native Kubernetes Cluster](#) or [Create a VMware Tanzu Kubernetes Grid Integrated Edition Cluster](#).

Add an Organization VDC Kubernetes Policy

You can add an organization VDC Kubernetes policy by using a provider VDC Kubernetes policy. Tenants 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. The 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 at least one flex organization VDC in your environment. See [Create an Organization Virtual Data Center](#).
- 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. For more information on vSphere with VMware Tanzu in VMware Cloud Director, see [Using Kubernetes with VMware Cloud Director](#).
- 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Organization VDCs**, and click the name of a flex organization VDC.

- 3 Under Policies, select **Kubernetes**, 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

[Manage Quotas on the Resource Consumption of an Organization](#)

Edit an Organization VDC Kubernetes Policy

You can modify an organization VDC Kubernetes policy to change its description and the CPU and memory limits.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Organization VDCs**, and click the name of a flex organization VDC.
- 3 Under Policies, select **Kubernetes**, select the policy you want to edit and click **Edit**.

The **Edit VDC Kubernetes Policy** wizard appears.

- 4 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.

- 5 Edit the CPU and Memory limit for the organization VDC Kubernetes policy and click **Next**.

You cannot edit the CPU and Memory reservation.

- 6 Review the new policy details and click **Save**.

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 [Using Kubernetes with VMware Cloud Director](#).

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

- Publish the Kubernetes Container Clusters plug-in to any organizations that you want to manage Tanzu Kubernetes 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](#).
- You must publish the **vmware:tkgcluster Entitlement** rights bundle to any organizations that you want to work with clusters. After sharing the rights bundle, you must add the **Edit: Tanzu Kubernetes Guest Cluster** right to the roles you want to create and modify Tanzu Kubernetes clusters. If you want the users also to delete clusters, you must add the **Full Control: Tanzu**

Kubernetes Guest Cluster right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).

- Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).

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 [Using Kubernetes with VMware Cloud Director](#).

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**.
- To enable the organization VDC for native Kubernetes cluster deployment, set up the Container Service Extension server. See the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- Publish the CSE native policy created during the CSE server setup to an organization VDC. To use the UI, see [Add a VM Placement Policy to an Organization VDC](#). Alternatively, you can use the CSE 3.0 CLI to publish policy by running the `vcd cse ovdc enable Organization_VDC_Name --org Organization_Name --native` command.
- You must publish the **cse:nativeCluster Entitlement** rights bundle to any organizations that you want to work with native clusters. After sharing the rights bundle, you must add the **Edit CSE:NATIVECLUSTER** right to the roles you want to create and modify Tanzu Kubernetes

clusters. If you want the users also to delete clusters, you must add the **Full Control CSE:NATIVECLUSTER** right to the roles. In addition, you can assign the administrator rights to users that you want to view all Tanzu Kubernetes clusters in an organization or users that you want to manage clusters across sites. For information about the rights and access levels for Runtime Defined Entities (RDEs), see [Chapter 14 Managing Defined Entities](#).

- Grant access to tenants or system administrators by creating Access Control List (ACL) entries. For more information on sharing Runtime Defined Entities (RDEs), see [Sharing Defined Entities](#).

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 [Using Kubernetes with VMware Cloud Director](#).

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

By using the TKGI enablement metadata, you can provide access to the tenants to create TKGI clusters and to access the TKGI-enabled organization VDC. If you want to limit the tenants' ability to create TKGI clusters, you can provide access only to the organization VDC. In this case, the tenants can manage existing TKGI clusters but cannot create new ones.

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**.
- To enable the organization VDC for TKGI Kubernetes cluster deployment, set up the Container Service Extension server. For information about using the CSE CLI to enable an organization VDC for TKGI, see the [CSE Server Management](#) chapter in the Container Service Extension (CSE) documentation.
- If you want to provide tenant access to TKGI creation and management, you must publish the **{cse}:PKS DEPLOY RIGHT** to the specific organizations, and add the **{cse}:PKS DEPLOY RIGHT** right to the roles you want to create and manage TKGI clusters. The **{cse}:PKS DEPLOY RIGHT** is created during the Container Service Extension server install.

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.

Create an Organization Virtual Data Center

To allocate resources to an organization, you must create an organization virtual data center (VDC). An organization virtual data center obtains its resources from a provider VDC. One organization can have multiple organization VDCs.

Prerequisites

Create a provider VDC. See [Create a Provider Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click **New**.
- 3 Enter a name and, optionally, a description for the new organization VDC.
- 4 (Optional) To deactivate the new organization VDC upon creation, turn off the **Enable the organization VDC** toggle.

Users cannot deploy vApps on a deactivated organization VDC.

- 5 Click **Next**.
- 6 Select the radio button next to the name of the organization to which you want to add this VDC, and click **Next**.
- 7 Select the radio button next to the name of the provider VDC from which you want the organization VDC to obtain compute and storage resources, and click **Next**.

The provider VDC list displays all activated provider VDC at the site with information about the available resources. The networks list displays information about the networks available to the selected provider VDC.

- 8 Select an allocation model for this organization VDC, and click **Next**.

Option	Description
Allocation pool	A percentage of the resources you allocate from the provider VDC are committed to the organization VDC. You can specify the percentage for both CPU and memory.
Pay-as-you-go	Resources are committed only when users create vApps in the organization VDC.

Option	Description
Reservation pool	All the resources you allocate are immediately committed to the organization VDC.
Flex	You can control the resource consumption at both the VDC and the individual virtual machine levels. The flex allocation model supports the capabilities of organization VDC compute policies. Flex allocation model supports all allocation configurations that are available in the other allocation models.

9 Configure the allocation settings for the allocation model that you selected, and click **Next**.

Option	Description	Allocation model
Elasticity	Activate or deactivate the elastic pool feature. An elastic organization VDC spans and uses all resource pools associated with its provider VDC.	Flex
Include VM Memory Overhead	Include or exclude memory overhead.	Flex
CPU allocation	The maximum amount of CPU that you want to allocate to the virtual machines running in this organization VDC.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Reservation Pool ■ Flex
Allow CPU resources to grow beyond	To provide unlimited CPU resources to this organization VDC, turn on this toggle.	Reservation Pool
CPU Quota	The maximum amount of CPU consumption for this organization VDC.	<ul style="list-style-type: none"> ■ Pay-as-you-go ■ Flex
CPU resources guaranteed	The percentage of CPU resources that you want to guarantee to a virtual machine running in this organization VDC. You can control overcommitment of CPU resources by guaranteeing less than 100 percent. For an Allocation Pool allocation model, the percentage guarantee also determines what percentage of the CPU allocation is committed for this organization VDC.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-as-you-go ■ Flex
vCPU Speed	The vCPU speed. Virtual machines running in the organization VDC are assigned this amount of GHz per vCPU.	<ul style="list-style-type: none"> ■ Pay-as-you-go ■ Flex
Memory allocation	The maximum amount of memory that you want to allocate to the virtual machines running in the organization VDC.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Reservation Pool
Memory Quota	The maximum amount of memory consumption for this organization VDC.	<ul style="list-style-type: none"> ■ Pay-as-you-go ■ Flex

Option	Description	Allocation model
Memory resources guaranteed	The percentage of memory resources that you want to guarantee to virtual machines running in the organization VDC. You can overcommit resources by guaranteeing less than 100 percent. For an Allocation Pool allocation model, the percentage guarantee also determines what percentage of the memory allocation is committed for this organization VDC.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-as-you-go ■ Flex
Maximum number of VMs	The maximum number of virtual machines that can exist in the organization VDC.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-as-you-go ■ Reservation Pool ■ Flex

10 Configure the storage settings for this organization VDC, and click **Next**.

The list contains the activated storage policies on the source provider VDC.

- a Select the check boxes of one or more storage policies that you want to add to this organization VDC.
- b (Optional) To limit the amount of the allocated storage capacity for a selected storage policy, select **Limited** from the drop-down menu in the **Allocation Type** cell, and enter the maximum capacity in the **Allocated Storage** cell.
- c (Optional) To change the default storage policy, from the **Default instantiation policy** drop-down menu, select the target default storage policy.

VMware Cloud Director uses the default storage policy for all virtual machine provisioning operations where the storage policy is not specified at the virtual machine or vApp template level.

- d (Optional) To activate thin provisioning for virtual machines in the organization VDC, turn on the **Thin provisioning** toggle.
- e (Optional) To deactivate fast provisioning for virtual machines in the organization VDC, turn off the **Fast provisioning** toggle.

11 Configure the network pool settings for this organization VDC, and click **Next**.

VMware Cloud Director uses the network pool to create vApp networks and internal organization VDC networks.

- To skip adding a network pool at this stage, turn off the **Use Network Pool** toggle.
- To configure a network pool, select the radio button next to the name of the target network pool, and enter the Quota for this organization VDC.

The quota is the maximum number of provisioned networks in the organization VDC backed by this network pool. Must not exceed the number of the available networks for the selected network pool.

Note Organization VDCs that are backed by NSX-T Data Center only support Geneve network pools.

- 12 Review the **Ready to Complete** page, and click **Finish**.

Activate or Deactivate an Organization Virtual Data Center

To prevent additional vApps and virtual machines from using compute and storage resources from an organization virtual data center, you can deactivate this organization virtual data center. Running vApps and powered on virtual machines continue to run, but you cannot create or start additional vApps or virtual machines.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Select the radio button next to the name of the target organization virtual data center, and click **Enable** or **Disable**.
- 4 To confirm, click **OK**.

Delete an Organization Virtual Data Center

To remove all resources of an organization virtual data center from an organization, you can delete this organization virtual data center. The resources remain unaffected in the source provider virtual data center.

Important This operation permanently removes the organization virtual data center and all its VMs, vApps, organization virtual data center networks, and edge gateways.

Prerequisites

If you want to keep certain VMs, vApps, vApp templates, or media files that belong to the target organization virtual data center, move them to another organization virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Select the radio button next to the name of the organization virtual data center that you want to remove, and click **Delete**.
- 4 If this organization virtual data center contains any resources, such as VMs, vApps, organization virtual data center networks, and edge gateways, to confirm their removal, select the check box for each resource type.
- 5 To confirm, click **Delete**.

Managing Virtual Data Center Templates

Starting with VMware Cloud Director 10.2.2, you can create and share virtual data center (VDC) templates with tenant organizations so that **organization administrators** can use the templates to create VDCs.

By creating and sharing VDC templates with organizations, you can enable self-service provisioning of organization VDCs while retaining administrative control over allocation of system resources such as provider VDCs and external networks.

A VDC template specifies the allocation model, memory, CPU resource configuration and storage policies for the new organization VDC, and optionally, an edge gateway and organization VDC network.

Create an Organization Virtual Data Center Template

Starting with VMware Cloud Director 10.2.2, you can use the HTML5 UI to create organization virtual data center (VDC) templates for VDCs backed by NSX Data Center for vSphere or NSX-T Data Center.

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.
- 2 In the left panel, select **Organization VDC Templates**, and click **New**.
- 3 Select a network provider type, select a provider VDC and external network pair, and click **Next**.

For NSX Data Center for vSphere, when a user instantiates an organization VDC from this template, VMware Cloud Director applies the selected edge clusters to the new organization VDC. All newly deployed edge gateways within the new organization VDC use these primary and secondary edge clusters as placements.

For NSX-T Data Center, VMware Cloud Director uses the **Services Edge Cluster** to deploy the networking services such as DHCP, VPN, and DNS services. VMware Cloud Director uses the **Edge Cluster For NSX-T Gateway** to deploy the gateway.

After you instantiate an organization VDC template, you cannot edit the edge clusters.

- 4 Select an allocation model for this organization VDC, and click **Next**.

Option	Description
Allocation pool	A percentage of the resources you allocate from the provider VDC are committed to the organization VDC. You can specify the percentage for both CPU and memory.
Pay-As-You-Go	Resources are committed only when users create vApps in the organization VDC.

Option	Description
Reservation pool	All the resources you allocate are immediately committed to the organization VDC.
Flex	You can control the resource consumption at both the VDC and the individual virtual machine levels. The flex allocation model supports the capabilities of organization VDC compute policies. Flex allocation model supports all allocation configurations that are available in the other allocation models.

5 Configure the allocation settings for the allocation model that you selected, and click **Next**.

Option	Description	Allocation Model
Elasticity	Activate or deactivate the elastic pool feature. An elastic organization VDC spans and uses all resource pools associated with its provider VDC.	Flex
Include VM Memory Overhead	Include or exclude memory overhead.	Flex
CPU allocation	The maximum amount of CPU that you want to allocate to the virtual machines running in this organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Reservation Pool ■ Flex
Allow CPU resources to grow beyond	To provide unlimited CPU resources to this organization virtual data center, turn on this toggle.	Reservation Pool
CPU Quota	The maximum amount of CPU consumption for this organization virtual data center.	<ul style="list-style-type: none"> ■ Pay-As-You-Go ■ Flex
CPU resources guaranteed	The percentage of CPU resources that you want to guarantee to a virtual machine running in this organization virtual data center. You can control overcommitment of CPU resources by guaranteeing less than 100 percent. For an Allocation Pool allocation model, the percentage guarantee also determines what percentage of the CPU allocation is committed for this organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-As-You-Go ■ Flex
vCPU Speed	The vCPU speed. Virtual machines running in the organization virtual data center are assigned this amount of GHz per vCPU.	<ul style="list-style-type: none"> ■ Pay-As-You-Go ■ Flex
Memory allocation	The maximum amount of memory that you want to allocate to the virtual machines running in the organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Reservation Pool
Memory Limit	The maximum amount of memory consumption for this organization virtual data center.	<ul style="list-style-type: none"> ■ Pay-As-You-Go ■ Flex

Option	Description	Allocation Model
Memory resources guaranteed	The percentage of memory resources that you want to guarantee to virtual machines running in the organization virtual data center. You can overcommit resources by guaranteeing less than 100 percent. For an Allocation Pool allocation model, the percentage guarantee also determines what percentage of the memory allocation is committed for this organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-As-You-Go ■ Flex
Maximum number of VMs	The maximum number of virtual machines that can exist in the organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-As-You-Go ■ Reservation Pool ■ Flex

6 Configure the storage settings for this organization virtual data center, and click **Next**.

The list contains the enabled storage policies on the source provider VDC.

- a Select one or more storage policies that you want to add to this organization VDC.
- b (Optional) To limit the amount of the allocated storage capacity for a selected storage policy, select **Limited** from the drop-down menu in the **Allocation Type** cell, and enter the maximum capacity in the **Allocated Storage** cell.
- c (Optional) To change the default storage policy, from the **Default instantiation policy** drop-down menu, select the target default storage policy.

VMware Cloud Director uses the default storage policy for all virtual machine provisioning operations where the storage policy is not specified at the virtual machine or vApp template level.

- d (Optional) To enable thin provisioning for virtual machines in the organization VDC, turn on the **Thin provisioning** toggle.
 - e (Optional) To deactivate fast provisioning for virtual machines in the organization VDC, turn off the **Fast provisioning** toggle.
- 7** (Optional) Create an edge gateway.

- a Enter a name, and optionally, a description of the new edge gateway.
- b If you are creating a template for a VDC backed by NSX Data Center for vSphere, you can customize the general edge gateway settings, and click **Next**.

General Setting	Description
Distributed Routing	Configures an advanced gateway to provide distributed logical routing.
FIPS Mode	Configures the edge gateway to use NSX FIPS mode.
High Availability	Enables automatic failover to a backup edge gateway.

- c If you are creating a template for a VDC backed by NSX Data Center for vSphere, you can change the edge gateway configuration for your system resources.

Configuration	Description
Compact	Requires less memory and fewer compute resources.
Large	Provides increased capacity and performance than the Compact configuration. Large and X-Large configurations provide identical security functions.
X-Large	Used for environments that have a load balancer with large numbers of concurrent sessions.
Quad Large	Used for high throughput environments. Requires a high connection rate.

- d (Optional) Specify how many IPs you want to allocate for the use of the gateway services.

8 Configure the organization VDC network, and click **Next**.

- a Enter a name, and optionally, a description of the network.
- b 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**.

- c To make the organization VDC network available to other organization VDCs within the same organization, turn on the **Shared** toggle.

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 model and run more VMs on that network on a temporary basis.

Note The Organization VDCs must share the same network pool.

9 Add and IP address range from the ranges of the available static IP pools, and click **Next**.

10 (Optional) Configure the network pool settings for the organization VDC, and click **Next**.

The quota is the maximum number of provisioned networks in the organization VDC backed by this network pool. The quota must not exceed the number of the available networks for the selected network pool.

11 Select the organizations that you want to view and instantiate VDCs from this template, and click **Next**.

System administrators can instantiate a VDC from any organization VDC template. By using the VMware Cloud Director Tenant Portal, **organization administrators** can instantiate a VDC if their organization is in the access list of a template.

12 Enter a system name and tenant-facing name of the template, and click **Next**.

13 Review the organization VDC template configuration and click **Finish**.

What to do next

- [Instantiate a Virtual Data Center from a Template.](#)
- [Edit an Organization VDC Template.](#) You can edit all the properties of an existing VDC template, except the network provider type.
- To create a copy of an organization VDC template that you can optionally customize, clone the template. The steps for cloning are similar to the steps for editing a template.
- Delete an organization VDC template.

Instantiate a Virtual Data Center from a Template

To create an organization virtual data center (VDC) from a VDC template, instantiate a VDC.

System administrators can instantiate a VDC from any organization VDC template. By using the VMware Cloud Director Tenant Portal, **organization administrators** can instantiate a VDC if their organization is in the access list of a template.

Prerequisites

[Create an Organization Virtual Data Center Template](#)

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.
- 2 In the left panel, select **Organization VDC Templates**.
- 3 Select an organization VDC template and click **Instantiate VDC**.
- 4 Enter a name, and optionally, a description of the new organization virtual data center.
- 5 Select an organization for the organization VDC and click **Create**.

Edit an Organization VDC Template

You can modify all the properties of an existing virtual data center (VDC) template, except the network provider type.

Prerequisites

[Create an Organization Virtual Data Center Template](#)

Procedure

- 1 From the top navigation bar, select **Resources**, and click **Cloud Resources**.
- 2 In the left panel, select **Organization VDC Templates**, and click **Edit**.

3 Select a provider VDC and external network pair, and click **Next**.

For NSX Data Center for vSphere, when a user instantiates an organization VDC from this template, VMware Cloud Director applies the selected edge clusters to the new organization VDC. All newly deployed edge gateways within the new organization VDC use these primary and secondary edge clusters as placements.

For NSX-T Data Center, VMware Cloud Director uses the **Services Edge Cluster** to deploy the networking services such as DHCP, VPN, and DNS services. VMware Cloud Director uses the **Edge Cluster For NSX-T Gateway** to deploy the gateway.

After you instantiate an organization VDC template, you cannot edit the edge clusters.

4 Select an allocation model for this organization VDC, and click **Next**.

Option	Description
Allocation pool	A percentage of the resources you allocate from the provider VDC are committed to the organization VDC. You can specify the percentage for both CPU and memory.
Pay-As-You-Go	Resources are committed only when users create vApps in the organization VDC.
Reservation pool	All the resources you allocate are immediately committed to the organization VDC.
Flex	You can control the resource consumption at both the VDC and the individual virtual machine levels. The flex allocation model supports the capabilities of organization VDC compute policies. Flex allocation model supports all allocation configurations that are available in the other allocation models.

5 Configure the allocation settings for the allocation model that you selected, and click **Next**.

Option	Description	Allocation Model
Elasticity	Activate or deactivate the elastic pool feature. An elastic organization VDC spans and uses all resource pools associated with its provider VDC.	Flex
Include VM Memory Overhead	Include or exclude memory overhead.	Flex
CPU allocation	The maximum amount of CPU that you want to allocate to the virtual machines running in this organization virtual data center.	<input type="checkbox"/> Allocation Pool <input type="checkbox"/> Reservation Pool <input type="checkbox"/> Flex
Allow CPU resources to grow beyond	To provide unlimited CPU resources to this organization virtual data center, turn on this toggle.	Reservation Pool
CPU Quota	The maximum amount of CPU consumption for this organization virtual data center.	<input type="checkbox"/> Pay-As-You-Go <input type="checkbox"/> Flex

Option	Description	Allocation Model
CPU resources guaranteed	The percentage of CPU resources that you want to guarantee to a virtual machine running in this organization virtual data center. You can control overcommitment of CPU resources by guaranteeing less than 100 percent. For an Allocation Pool allocation model, the percentage guarantee also determines what percentage of the CPU allocation is committed for this organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-As-You-Go ■ Flex
vCPU Speed	The vCPU speed. Virtual machines running in the organization virtual data center are assigned this amount of GHz per vCPU.	<ul style="list-style-type: none"> ■ Pay-As-You-Go ■ Flex
Memory allocation	The maximum amount of memory that you want to allocate to the virtual machines running in the organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Reservation Pool
Memory Limit	The maximum amount of memory consumption for this organization virtual data center.	<ul style="list-style-type: none"> ■ Pay-As-You-Go ■ Flex
Memory resources guaranteed	The percentage of memory resources that you want to guarantee to virtual machines running in the organization virtual data center. You can overcommit resources by guaranteeing less than 100 percent. For an Allocation Pool allocation model, the percentage guarantee also determines what percentage of the memory allocation is committed for this organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-As-You-Go ■ Flex
Maximum number of VMs	The maximum number of virtual machines that can exist in the organization virtual data center.	<ul style="list-style-type: none"> ■ Allocation Pool ■ Pay-As-You-Go ■ Reservation Pool ■ Flex

6 Configure the storage settings for this organization virtual data center, and click **Next**.

The list contains the enabled storage policies on the source provider VDC.

- a Select one or more storage policies that you want to add to this organization VDC.
- b (Optional) To limit the amount of the allocated storage capacity for a selected storage policy, select **Limited** from the drop-down menu in the **Allocation Type** cell, and enter the maximum capacity in the **Allocated Storage** cell.
- c (Optional) To change the default storage policy, from the **Default instantiation policy** drop-down menu, select the target default storage policy.

VMware Cloud Director uses the default storage policy for all virtual machine provisioning operations where the storage policy is not specified at the virtual machine or vApp template level.
- d (Optional) To enable thin provisioning for virtual machines in the organization VDC, turn on the **Thin provisioning** toggle.
- e (Optional) To deactivate fast provisioning for virtual machines in the organization VDC, turn off the **Fast provisioning** toggle.

7 (Optional) Create an edge gateway.

- a Enter a name, and optionally, a description of the new edge gateway.
- b If you are editing a template for a VDC backed by NSX Data Center for vSphere, you can customize the general edge gateway settings, and click **Next**.

General Setting	Description
Distributed Routing	Configures an advanced gateway to provide distributed logical routing.
FIPS Mode	Configures the edge gateway to use NSX FIPS mode.
High Availability	Enables automatic failover to a backup edge gateway.

- c If you are editing a template for a VDC backed by NSX Data Center for vSphere, you can change the edge gateway configuration for your system resources.

Configuration	Description
Compact	Requires less memory and fewer compute resources.
Large	Provides increased capacity and performance than the Compact configuration. Large and X-Large configurations provide identical security functions.
X-Large	Used for environments that have a load balancer with large numbers of concurrent sessions.
Quad Large	Used for high throughput environments. Requires a high connection rate.

- d (Optional) Specify how many IPs you want to allocate for the use of the gateway services.

8 Configure the organization VDC network, and click **Next**.

- a Enter a name, and optionally, a description of the network.
- b 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**.

- c To make the organization VDC network available to other organization VDCs within the same organization, turn on the **Shared** toggle.

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 model and run more VMs on that network on a temporary basis.

Note The Organization VDCs must share the same network pool.

9 Add and IP address range from the ranges of the available static IP pools, and click **Next**.

- 10 (Optional) Configure the network pool settings for the organization VDC, and click **Next**.

The quota is the maximum number of provisioned networks in the organization VDC backed by this network pool. The quota must not exceed the number of the available networks for the selected network pool.

- 11 Select the organizations that you want to view and instantiate VDCs from this template, and click **Next**.
- 12 Enter a system name and tenant-facing name of the template, and click **Next**.
- 13 Review the organization VDC template configuration and click **Finish**.

Modify the Name and the Description of an Organization Virtual Data Center

As your VMware Cloud Director installation expands, you might want to assign a more meaningful name or description to an existing organization virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 On the **General** tab, in the upper right corner, click **Edit**.
- 4 Enter a new name and description, and click **Save**.

Modify the Allocation Model Settings of an Organization Virtual Data Center

You cannot change the allocation model for an organization virtual data center, but you can change the allocation settings for the allocation model that you specified during the creation of the organization virtual data center.

You can modify the allocation settings for the allocation model that you configured during the creation of the organization virtual data center. See [Step 9](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 On the **Allocation** tab, in the upper right corner, click **Edit**.
- 4 Edit the allocation model settings, and click **Save**.

Modifying the Storage Settings of an Organization Virtual Data Center

You can modify the storage settings that you configured during the creation of the organization virtual data center.

Enabling VM Encryption on Storage Policies of an Organization Virtual Data Center

You can add an encryption-enabled storage policy to an organization VDC. You can encrypt VMs and disks by associating a VM or disk with a storage policy that has the VM Encryption capability.

Starting with VMware Cloud Director 10.1, you can improve the security of your data by using VM encryption. 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 provider VDC has a storage policy with enabled VM Encryption, you can add the encryption-enabled policy to an organization VDC. See [Enabling VM Encryption on Storage Policies of a Provider Virtual Data Center](#) and [Add a VM Storage Policy to an Organization Virtual Data Center](#). After that, by using the VMware Cloud Director Tenant Portal, tenants can associate a VM or disk with a storage policy with enabled VM Encryption.

VM Encryption Limitations

The following actions are not supported in VMware Cloud Director 10.1.

- 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 VMs and disks. For more information about roles and rights, see [Predefined Roles and Their Rights](#).

You can view all storage capabilities in the **Capabilities** column under **Resources > vSphere Resources > Storage Policies**. This column displays the VM encryption, tag-based association, vSAN, and IOPS limiting storage capabilities. To view the full list of storage capabilities, expand the row by clicking the arrow on the left side of the storage policy name.

You can also view the storage capability information in the **Storage** tab of an organization VDC.

Modify the VM Provisioning Settings of an Organization Virtual Data Center

You can modify the virtual machine thin provisioning and fast provisioning settings that you configured during the creation of the organization virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage** and click **Edit**.
- 4 (Optional) Modify the thin provisioning setting.
 - To deactivate thin provisioning for virtual machines in the organization virtual data center, turn off the **Thin provisioning** toggle.
 - To activate thin provisioning for virtual machines in the organization virtual data center, turn on the **Thin provisioning** toggle.
- 5 (Optional) Modify the fast provisioning setting.
 - To activate fast provisioning for virtual machines in the organization virtual data center, turn on the **Fast provisioning** toggle.
 - To deactivate fast provisioning for virtual machines in the organization virtual data center, turn off the **Fast provisioning** toggle.
- 6 Click **Edit**.

Add a VM Storage Policy to an Organization Virtual Data Center

You can configure an organization virtual data center to support a VM storage policy that you previously added to the backing provider virtual data center.

Prerequisites

You added the target VM storage policy to the source provider virtual data center. See [Add a VM Storage Policy to a Provider Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage** , and click **Add**.

You can see a list of the available additional storage policies in the source provider virtual data center.

- 4 Select the check boxes of one or more storage policies that you want to add, and click **Add**.

Change the Default Storage Policy on an Organization Virtual Data Center

You can change the default storage policy that you configured during the creation of an organization virtual data center.

VMware Cloud Director uses the default storage policy for all virtual machine provisioning operations where the storage policy is not specified at the virtual machine or vApp template level.

Prerequisites

- The target default storage policy is added to the organization virtual data center. See [Add a VM Storage Policy to an Organization Virtual Data Center](#).
- The target default storage policy is enabled on the organization virtual data center. See [Activate or Deactivate a Storage Policy on an Organization Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the name of the target default storage policy, and click **Set as default**.
- 5 To confirm, click **OK**.

Edit the Limit of a Storage Policy on an Organization Virtual Data Center

You can change the limit of the allocated storage capacity that you configured for a storage policy during the creation of an organization virtual data center.

You can set the allocated storage capacity as unlimited or configure a maximum amount of allocated storage capacity for a storage policy on an organization virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage**.
- 4 Click the radio button next to the name of the target storage policy, and click **Edit limit**.
- 5 Configure the limit setting for this storage policy.
 - To set a limit, select the upper radio button, and enter the maximum amount of storage resource for this storage policy on this organization virtual data center.
 - To set no limit, select the **Unlimited** radio button.
- 6 Click **Edit**.

Modify the Metadata for a VM Storage Policy on an Organization Virtual Data Center

You can add, edit, and delete metadata for a storage policy on an organization virtual data center.

By using object metadata, you can associate user-defined *name=value* pairs with a storage policy on an organization virtual data center. You can use object metadata in vCloud API query filter expressions.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage**.
- 4 Click the radio button next to the name of the target storage policy, and click **Metadata**.
- 5 Click **Edit**.
- 6 (Optional) To add a key-value pair, click **Add**, enter a name and a value, and select a type for the new key-value pair.

- 7 (Optional) To edit a key-value pair, enter a new name and a value, and select a new type for the key-value pair.
- 8 (Optional) To remove a key-value pair, in the right end of the row, click the **Delete** icon.
- 9 Click **Save**, and click **OK**.

Activate or Deactivate a Storage Policy on an Organization Virtual Data Center

To prevent additional vApps and virtual machines from using a storage policy on an organization virtual data center, you can deactivate this storage policy on the organization virtual data center. Running vApps and powered on virtual machines continue to run, but you cannot create or start additional vApps or virtual machines on this storage policy.

You cannot deactivate the default storage policy.

Prerequisites

If you want to deactivate the default storage policy, [Change the Default Storage Policy on an Organization Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage**.
- 4 Click the radio button next to the name of the target storage policy, and click **Enable** or **Disable**.
- 5 To confirm, click **OK**.

Delete a Storage Policy from an Organization Virtual Data Center

To prevent an organization virtual data center from using a storage policy, you can remove this storage policy from the organization virtual data center. Running vApps and powered on virtual machines continue to run, but you cannot create or start additional vApps or virtual machines on this storage policy.

Prerequisites

Deactivate the storage policy that you want to remove. See [Activate or Deactivate a Storage Policy on an Organization Virtual Data Center](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.

- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the name of the target storage policy, and click **Remove**.
- 5 To confirm, click **Remove**.

Edit the Organization VDC Storage Policy Settings

You can change the I/O operations per second (IOPS) settings of an organization VDC storage policy. By default, the organization VDC storage policies inherit the provider VDC storage policy settings. You can customize the settings per organization VDC storage policy.

Prerequisites

[Add a VM Storage Policy to an Organization Virtual Data Center](#)

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Under **Policies**, select **Storage** .
- 4 Click the radio button next to the target storage policy, and click **Edit Settings**.
- 5 If you want the IOPS settings of the organization VDC storage policy to be different from the provider VDC storage policy, turn off the **Inherit From Provider VDC** toggle.
- 6 If you want to limit the I/O operations per second, turn on the **IOPS Limiting Enabled** toggle.
- 7 If you want IOPS to be considered during placement, turn on the **Impact Placement** toggle.

If the **Impact Placement** toggle is turned on, VMware Cloud Director provides IOPS load balancing across datastores. When you set IOPS settings for a disk, VMware Cloud Director considers datastores with enough IOPS capacity for the selected disk. If the **Impact Placement** toggle is turned off, you do not need to set IOPS capacities per datastore and you can use Storage DRS clusters.

- 8 (Optional) Configure the maximum and default IOPS settings.
- 9 Click **Save**.

Edit the Network Settings of an Organization Virtual Data Center

You can change the network pool from which new networks are provisioned in an organization virtual data center. You can also enable organization virtual data centers to become eligible for cross-virtual data center networking.

A network pool is a group of undifferentiated networks that you can use to create vApp networks, routed organization VDC networks, and internal organization VDC networks. You can change the network pool for new networks. Existing networks continue to use the old network pools.

With organization virtual data centers that are enabled for cross-virtual data center networking, organization users with relevant rights can create data center groups and stretched layer 2 networks in these groups.

Prerequisites

If you want to enable cross-VDC networking for an organization virtual data center, verify that you configured cross-vCenter NSX on the backing provider virtual data center.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 On the **Network Pool** tab, in the upper right corner, click **Edit**.
You can see the number of the networks used by this organization virtual data center.
- 4 (Optional) Configure the network pool settings for this organization virtual data center.

Note Organization VDCs that are backed by NSX-T Data Center only support Geneve network pools.

- If you do not want a network pool for this organization virtual data center, turn off the **Use network pool** toggle.
 - If you want to configure a network pool for this organization virtual data center, follow these steps:
 - a Turn on the **Use network pool** toggle.
You can see a list of the available network pools with information about their use, available networks, and capacity.
 - b Select the radio button next to the name of the target resource pool.
 - c Configure the quota for this network pool in this organization virtual data center.
The quota is the maximum number of provisioned networks. Must not exceed the number of the available networks for the selected network pool.
- 5 To enable cross-virtual data center networking for this organization virtual data center, turn on the **Cross VDC Networking** toggle.
 - 6 Click **Save**.

Results

In the VMware Cloud Director Tenant Portal, the virtual data centers that enabled for cross-virtual data center networking appear in the list of data centers for creating a data center group. For information about creating data center groups, see the *VMware Cloud Director Tenant Portal Guide*.

Configuring Cross-Virtual Data Center Networking

The cross-virtual data center networking feature enables organizations that have virtual data centers backed by multiple vCenter Server instances to stretch layer 2 networks across up to four virtual data centers. Cross-virtual data center networking relies on cross-vCenter NSX and can span multiple VMware Cloud Director sites.

Cross-virtual data center networking requires NSX Data Center for vSphere.

With cross-virtual data center networking, organizations can group up to four virtual data centers and configure egresses and stretched layer 2 networks in each group.

The participating organization virtual data centers can belong to different VMware Cloud Director sites. See [Configuring and Managing Multisite Deployments](#).

Organizations can use cross-virtual data center networking to implement high availability solutions or distributed systems architectures, where an application can be distributed across multiple virtual data centers or sites.

The **system administrator** must configure the underlying cross-vCenter NSX environment, the VMware Cloud Director servers, and enable cross-virtual data center networking for each virtual data center.

- 1 Configure one of the NSX Manager instances as a Primary NSX Manager instance. See the *Cross-vCenter NSX Installation Guide*.
 - a Deploy the NSX cluster on the primary NSX Manager instance.
 - b Prepare the ESXi hosts on the primary NSX Manager instance.
 - c Configure VXLAN from the primary NSX Manager instance.
 - d Assign the primary role to the NSX Manager instance.
 - e Create a pool for segment IP for the universal transport zone.
 - f Add a universal transport zone.
- 2 Configure the rest of the NSX Manager instances as Secondary NSX Managers. See the *Cross-vCenter NSX Installation Guide*.
 - a Prepare the ESXi hosts on each secondary NSX Manager instance.
 - b Configure VXLAN from each secondary NSX Manager instance.
 - c Assign the secondary role to each NSX Manager instance.
 - d Connect the ESXi clusters to the universal transport zone.

- 3 Configure the control VM properties for each NSX Manager instance. See [Modify NSX Manager Settings](#).
- 4 Create a VXLAN backed network pool using a universal type transport zone from any vCenter Server instance. See [Create a Network Pool Backed by an NSX Data Center for vSphere Transport Zone](#).

Note For multisite deployments, you must create a VXLAN backed network pool in each VMware Cloud Director site.

- 5 Enable cross-virtual data center networking on each organization virtual data center. See [Edit the Network Settings of an Organization Virtual Data Center](#).
- 6 If the organization has multisite virtual data centers, verify that the installation IDs in the different VMware Cloud Director sites are different. If there are VMware Cloud Director sites that are configured with the same installation ID, see [Regenerating MAC Addresses for Multisite Stretched Networks](#) in the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

The **organization administrator** can now create and configure data center groups, egresses, and stretched networks. For information about managing cross-virtual data center networking, see the *VMware Cloud Director Tenant Portal Guide*.

Modify the Metadata for an Organization Virtual Data Center

You can add, edit, and delete metadata for an organization virtual data center.

By using object metadata, you can associate user-defined `name=value` pairs with an organization virtual data center. You can use object metadata in vCloud API query filter expressions.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Metadata** tab.
- 4 Click **Edit**.
- 5 (Optional) To add a key-value pair, click **Add**, enter a name and a value, and select a type for the new key-value pair.
- 6 (Optional) To edit a key-value pair, enter a new name and a value, and select a new type for the key-value pair.
- 7 (Optional) To remove a key-value pair, in the right end of the row, click the **Delete** icon.
- 8 Click **Save**, and click **OK**.

View the Resource Pools of an Organization Virtual Data Center

You can view a list of the vCenter Server resource pools that an organization virtual data center uses.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Resource Pools** tab.

Results

You can see a table with the resource pools in use by the organization virtual data center, and the vCenter Server instance to which each resource pool belongs.

Managing the Distributed Firewall on an Organization Virtual Data Center

To provide Layer 3 and Layer 2 network security in an organization virtual data center, you can enable and create rules for the distributed firewall on this organization virtual data center. With the distributed firewall rules, you can protect traffic traveling between virtual machines in an organization virtual data center.

VMware Cloud Director supports distributed firewall services on organization virtual data centers that are backed by NSX Data Center for vSphere.

For creating the distributed firewall rules, you can use various grouping objects and security groups. See [Custom Grouping Objects](#) and [Working with Security Groups](#).

For information about protecting traffic to and from an edge gateway, see [Managing an NSX Data Center for vSphere Edge Gateway Firewall](#).

Activate the Distributed Firewall on an Organization Virtual Data Center

Before you can manage the distributed firewall settings on an organization virtual data center, you must activate the distributed firewall on this organization virtual data center.

VMware Cloud Director supports distributed firewall services on organization virtual data centers that are backed by NSX Data Center for vSphere.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.

- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 On the **Distributed Firewall > General** tab, turn on the **Enable Distributed firewall** toggle.

Results

You can see the default firewall rules, which allow all Layer 3 and Layer 2 traffic to pass through the organization virtual data center.

- On the **Distributed Firewall > General** tab, you can see the default distributed firewall rule for Layer 3 traffic, named Default Allow Rule.
- On the **Distributed Firewall > Ethernet** tab, you can see the default distributed firewall rule for Layer 2 traffic, named Default Allow Rule.

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

- [Activate the Distributed Firewall on an Organization Virtual Data Center](#)
- 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 an 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 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.

- 5 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.

- 6 Click in the **Name** cell and type in a name.
- 7 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 General tab. Type the source value you want to use. Valid values are an IP address, CIDR, an IP range, or the keyword any . 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"> ■ Use the Select objects window to add objects that match your selections and click Keep to add them to the rule. ■ To exclude a source from the rule, add it to this rule using the Select objects window and then select the toggle exclusion icon to exclude that source from this rule. <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 Select objects window</p>

- 8 Click in the **Destination** cell and perform one of the following actions:

Action	Description
Click the IP icon	Applicable for rules defined on the General tab. Type the destination value you want to use. Valid values are an IP address, CIDR, an IP range, or the keyword any . 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"> ■ Use the Select objects window to add objects that match your selections and click Keep to add them to the rule. ■ To exclude a source from the rule, add it to this rule using the Select objects window and then select the toggle exclusion icon to exclude that source from this rule. <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 Select objects window</p>

- 9 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"> Select the service protocol. Type the port numbers for the source and destination ports, or specify any, and click Keep.
Click the + icon	To select a pre-defined service or service group, or define a new one: <ol style="list-style-type: none"> Select one or more objects and add them to the filter. Click Keep.

- 10 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.

- 11 In the **Direction** cell of the new rule, select whether the rule applies to incoming traffic, outgoing traffic, or both.
- 12 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**.
- 13 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.

- 14 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

- From the top navigation bar, select **Resources** and click **Cloud Resources**.
- In the left panel, click **Organization VDCs**.

- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 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.
- 5 Click **Save Changes**.

Custom Grouping Objects

The NSX 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. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Organization VDCs. c Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. d Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Edge Gateways. c Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. d Click the Grouping Objects tab.

- 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 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 an MAC set by using the **Grouping Objects** page. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ol style="list-style-type: none"> From the top navigation bar, under Resources, select Cloud Resources. In the left panel, click Organization VDCs. Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ol style="list-style-type: none"> From the top navigation bar, under Resources, select Cloud Resources. In the left panel, click Edge Gateways. Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. Click the Grouping Objects tab.

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 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 by using the **Grouping Objects** page. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ul style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Organization VDCs. c Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. d Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ul style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Edge Gateways. c Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. d Click the Grouping Objects tab.

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 by using the **Grouping Objects** page. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ul style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Organization VDCs. c Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. d Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ul style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Edge Gateways. c Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. d Click the Grouping Objects tab.

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.

Working with Security Groups

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 Click the **Grouping Objects > Security Groups** tab.

- 5 Click the **Create** () button.

- 6 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.

- 7 (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**.
- 8** (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.
- 9** (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.
- 10** 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** From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2** In the left panel, click **Organization VDCs**.
- 3** Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4** Click the **Grouping Objects > Security Groups** tab.
- 5** Select the security group you want to edit.
The details for the security group display below the list of security groups.
- 6** (Optional) Edit the name and the description of the security group.
- 7** (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**.
- 8 (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**.
 - 9 (Optional) Delete a dynamic member set by clicking the **Delete** icon next to the member set that you want to delete.
 - 10 (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.
 - 11 (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.
 - 12 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 Click the **Grouping Objects > Security Groups** tab.

- 5 Select the security group you want to delete.
- 6 Click the **Delete** button.
- 7 To confirm the deletion, click **OK**.

Results

The security group is deleted.

Working with Security Tags

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 Click the **Security Tags** tab.
- 5 Click the **Create** () button, and enter a name for the security tag.
- 6 (Optional) Enter a description for the security tag.

- 7 (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.

- 8 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 Click the **Security Tags** tab.
- 5 From the list of security tags, select the security tag that you want to edit, and click the **Edit**  button..
- 6 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.
- 7 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.
- 8 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 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**.
- 5 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 Click the **Security Tags** tab.
- 5 From the list of security tags, select the security tag that you want to edit.
- 6 Click the **Edit** () button.
- 7 Edit the name and the description of the security tag.
- 8 Assign the tag to or remove the assignment from the virtual machines that you select.
- 9 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](#)

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**.
- 3 Click the radio button next to the target organization virtual data center, and click **Manage Firewall**.
- 4 Click the **Security Tags** tab.
- 5 From the list of security tags, select the security tag that you want to delete.

6 Click the **Delete** () button.

7 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](#).

Managing NSX Data Center for vSphere Edge Gateways

7

An NSX Data Center for vSphere edge gateway provides a routed organization virtual data center 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.

Starting with VMware Cloud Director 9.7, the compute workload and the networking workload are isolated by using different vSphere resource pools and storage policies. Edge gateways reside on edge clusters that you must previously create. See [Working with NSX Data Center for vSphere Edge Clusters](#).

You can migrate legacy edge gateways to the corresponding edge clusters by redeploying these edge gateways. See [Redeploy an Edge Gateway](#).

Important Starting with version 9.7, VMware Cloud Director supports only advanced edge gateways. You must convert any legacy non-advanced edge gateway to an advanced gateway. See <https://kb.vmware.com/kb/66767>.

This chapter includes the following topics:

- [Working with NSX Data Center for vSphere Edge Clusters](#)
- [Add an NSX Data Center for vSphere Edge Gateway](#)
- [Configuring NSX Data Center for vSphere Edge Gateway Services](#)
- [View the Networks Use and IP Allocations on an Edge Gateway](#)
- [Editing Edge Gateway Properties](#)
- [Redeploy an Edge Gateway](#)
- [Delete an Edge Gateway](#)
- [Statistics and Logs for an Edge Gateway](#)
- [Enable SSH Command-Line Access to an Edge Gateway](#)

Working with NSX Data Center for vSphere Edge Clusters

To isolate the compute workloads from the networking workloads, VMware Cloud Director supports the edge cluster object. An edge cluster consists of a vSphere resource pool and

a storage policy that are used only for organization VDC edge gateways. Provider virtual data centers cannot use resources dedicated to edge clusters, and edge clusters cannot use resources dedicated to provider virtual data centers.

Edge clusters provide a dedicated L2 broadcast domain, which reduces the VLAN sprawls and ensures the network security and isolation. For example, the edge cluster can contain additional VLANs for peering with physical routers.

You can create any number of edge clusters. You can assign an edge cluster to an organization VDC as a primary or secondary edge cluster.

- The primary edge cluster for an organization VDC is used for the main edge appliance of an organization VDC edge gateway.
- The secondary edge cluster for an organization VDC is used for the standby edge appliance when an edge gateway is in HA mode.

Different organization VDCs can share edge clusters or can have their own dedicated edge clusters.

Starting with vCloud Director 9.7, the old process for using metadata to control the edge gateway placement is deprecated. See <https://kb.vmware.com/kb/2151398>.

You can migrate legacy edge gateways to newly created edge clusters by redeploying these edge gateways. See [Redeploy an Edge Gateway](#).

Preparing Your Environment for an Edge Cluster

- 1 In vSphere, create the resource pool for the target edge cluster.

If an organization virtual data center is using a VLAN network pool, the VLAN network pool and the edge cluster for this organization virtual data center must reside on the same vSphere distributed switch.

- 2 If an organization virtual data center is using a VXLAN network pool, in NSX, add the edge cluster to the VXLAN transport zone, after which synchronize the VXLAN network pool in VMware Cloud Director.

- 3 In vSphere, create the edge cluster storage profile.

Creating and Managing Edge Clusters

After you prepare your environment, to create and manage edge clusters, you must use the VMware Cloud Director OpenAPI `EdgeClusters` methods. See *Getting Started with VMware Cloud Director OpenAPI* at <https://code.vmware.com>.

Viewing edge clusters requires the **Edge Cluster View** right. Creating, updating, and deleting edge clusters require the **Edge Cluster Manage** right.

When you create an edge cluster, you specify the name, the vSphere resource pool, and the storage profile name.

After you create an edge cluster, you can modify its name and description. After you delete or move its containing edge gateways, you can delete an edge cluster.

Assigning an Edge Cluster to an Organization VDC

After you create an edge cluster, you can assign this edge cluster to an organization VDC by updating the organization VDC network profile. You can assign an edge cluster to an organization VDC as a primary or secondary edge cluster.

If you do not assign a secondary edge cluster, the standby edge appliance of an edge gateway in HA mode is deployed on the primary edge cluster but on a host different from the host running the primary edge appliance.

To update, view, and delete organization VDC network profiles, you must use the VMware Cloud Director OpenAPI `VdcNetworkProfile` methods. See *Getting Started with VMware Cloud Director OpenAPI* at <https://code.vmware.com>.

Considerations:

- The primary and secondary edge clusters must reside on the same vSphere distributed switch.
- If the organization VDC uses a VXLAN network pool, the NSX Transport Zone must span the compute and the edge clusters.
- If the organization VDC uses a VLAN network pool, the edge clusters and the compute clusters must be on the same vSphere distributed switch.

If you update again the primary or secondary edge cluster of an organization VDC, to move an existing edge gateway to the new cluster, you must redeploy this edge gateway. See [Redeploy an Edge Gateway](#).

Add an NSX Data Center for vSphere Edge Gateway

An NSX Data Center for vSphere 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.

Starting with VMware Cloud Director 9.7, NSX Data Center for vSphere edge gateways are deployed on edge clusters that you previously created and assigned to the organization VDC.

You can add an IPv4 or IPv6 edge gateway that connects to one or more external networks.

Note IPv6 edge gateways support limited services. IPv6 edge gateways support edge firewalls, distributed firewalls, and static routing.

Prerequisites

- For information about the system requirements for deploying an NSX Data Center for vSphere edge gateway, see the *NSX Administration Guide*.

- If you want to deploy the edge gateway on a dedicated edge cluster, create and assign an edge cluster to the organization virtual data center. See [Working with NSX Data Center for vSphere Edge Clusters](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left pane, click **Edge Gateways** and click **New**.
- 3 Select the NSX-V backed organization virtual data center on which you want to create the edge gateway, and click **Next**.
- 4 Enter a name and, optionally, a description for the new edge gateway.
- 5 Turn on or leave turned off each of these general edge gateway settings.

General Setting	Description
Distributed Routing	Configures the edge gateway to provide distributed logical routing.
FIPS Mode	Configures the edge gateway to use NSX FIPS mode.
High Availability	Enables automatic failover to a backup edge gateway.

- 6 Select the edge gateway configuration for your system resources and click **Next**.

Configuration	Description
Compact	Requires less memory and fewer compute resources.
Large	Provides increased capacity and performance than the Compact configuration. Large and X-Large configurations provide identical security functions.
X-Large	Used for environments that have a load balancer with large numbers of concurrent sessions.
Quad Large	Used for high throughput environments. Requires a high connection rate.

- 7 Select one or more subnets from the external networks to which the edge gateway can connect, and click **Next**.

If you assigned an edge cluster to the organization VDC, the displayed list contains the external networks that are accessible to this edge cluster.

- 8 (Optional) Configure a network as the default gateway.
 - a Turn on the **Configure default gateway** toggle.
 - b Click the radio button next to the name of the target external network, and click the radio button next to the target IP address.
 - c (Optional) Turn on the **Use default gateway for DNS Relay** toggle.
- 9 Click **Next**.

- 10 Turn on or leave turned off each of these advanced edge gateway settings, and click **Next**.

Advanced Setting	Description
IP Settings	You can manually enter an IP address for each subnet on the edge gateway.
Sub-Allocate IP Pools	You can suballocate multiple static IP pools from the available IP pools of each external network on the edge gateway.
Rate Limits	You can configure the inbound and outbound rate limits for each external network on the edge gateway.

- 11 (Optional) If you enabled one or more advanced settings in [Step 10](#), configure each enabled setting.

Advanced Setting	Steps
IP Settings	<p>For each network on the edge gateway, in the IP Addresses cell, enter an IP address, and click Next.</p> <p>If you do not enter an IP address for a network, the system assigns an arbitrary IP address to this network.</p>
Sub-Allocate IP Pools	<ol style="list-style-type: none"> 1 Click the radio button next to the name of an external network and click Edit. You can see the available IP pools for this external network and the current suballocated IP pools, if configured. 2 Edit the suballocated IP pools for this external network and click Save. You can add IP addresses and ranges from the ranges of the available IP pools. 3 Click Save. The system combines overlapping IP ranges. 4 Click Next. <p>Note Allocating IP addresses to an edge gateway is a process where the provider assigns ownership of IP addresses to the gateway. VMware Cloud Director automatically configures the appropriate gateway interface with the secondary addresses during the allocation process. If any of the IP addresses are used outside of VMware Cloud Director, this can cause IP address conflicts.</p>
Rate limits	For each external network on the edge gateway, turn on the Enable toggle, enter the limits in the Incoming Rate and Outgoing Rates cells, and click Next .

- 12 Review the **Ready to Complete** page, and click **Finish**.

Configuring NSX Data Center for vSphere Edge Gateway Services

You can configure services such as DHCP, firewall, network address translation (NAT), and VPN on an edge gateway.

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 the Distributed Firewall on an Organization Virtual Data Center](#).

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.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 any . 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"> ■ Use the Select objects window to add objects that match your selections and click Keep to add them to the rule. ■ To exclude a source from the rule, add it to this rule using the Select objects window and then select the toggle exclusion icon to exclude that source from this rule. <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 Select objects 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 any . 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"> ■ Use the Select objects window to add objects that match your selections and click Keep to add them to the rule. ■ To exclude a source from the rule, add it to this rule using the Select objects window and then select the toggle exclusion icon to exclude that source from this rule. <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 Select objects 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
Accept	Allows traffic from or to the specified sources, destinations, and services.
Deny	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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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**.

Apply Syslog Server Settings to an NSX Data Center for vSphere Edge Gateway

If you enabled logging for one or more edge gateway firewall rules, the edge gateway connects to the syslog server. If you created an edge gateway before the initial configuration of the syslog

server, or if you changed the syslog server settings, you must synchronize the syslog server settings for this edge gateway.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the radio button next to the name of the target edge gateway, and click **Sync syslog**.
- 4 To confirm, click **OK**.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 Primary Name Server and Secondary Name Server are set to Auto .
Primary Name Server	When you do not enable Auto Configure DNS , type your primary DNS server IP address of your primary DNS server. This IP address is used for hostname-to-IP address resolution.
Secondary Name Server	When you do not enable Auto Configure DNS , type your secondary DNS server IP address. This IP address is used for hostname-to-IP address resolution.
Default Gateway	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.
Subnet Mask	Type the subnet mask of the edge gateway interface.

Option	Description
Lease Never Expires	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, Lease Time is set to infinite.
Lease Time (Seconds)	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**.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, and click **Services**.
- 2 On the **DHCP > Bindings** tab, click the **Create** () button, specify the details for the binding, and click **Keep**.

Option	Description
MAC Address	Type the MAC address of the virtual machine that you want bound to the IP address.
Host Name	Type the host name you want set for that virtual machine when the virtual machine requests a DHCP lease.
IP Address	Type the IP address you want bound to the MAC address.
Subnet Mask	Type the subnet mask of the edge gateway interface.
Domain Name	Type the domain name of the DNS server.

Option	Description
Auto Configure DNS	Enable this toggle to use the DNS service configuration for this DNS binding. If enabled, the Primary Name Server and Secondary Name Server are set to Auto .
Primary Name Server	When you do not select Auto Configure DNS , type your primary DNS server IP address of your primary DNS server. This IP address is used for hostname-to-IP address resolution.
Secondary Name Server	When you do not select Auto Configure DNS , type your secondary DNS server IP address. This IP address is used for hostname-to-IP address resolution.
Default Gateway	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.
Lease Never Expires	Enable this toggle to keep the IP address bound to that MAC address forever. When you select this option, Lease Time is set to infinite.
Lease Time (Seconds)	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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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**.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Applied On	Select the interface on which to apply the rule.
Original IP/Range	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.
Protocol	Select the protocol to which the rule applies. To apply this rule on all protocols, select Any .
Original Port	(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 Protocol is set to ICMP or Any .
ICMP Type	When you select ICMP (an error reporting and a diagnostic utility used between devices to communicate error information) for Protocol , select the ICMP Type from the drop-down menu. ICMP messages are identified by the type field. By default, the ICMP type is set to any.
Translated IP/Range	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.
Translated Port	(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.
Source IP address	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.
Source Port	(Optional) Enter a port number for the source.
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

You can configure the static and dynamic routing capabilities that are provided by the NSX software for 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 provides, see *Routing* in the *NSX Administration* 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 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Network	Type the network in CIDR notation.
Next Hop	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.
MTU	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.
Interface	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.
Description	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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Enable Graceful Restart	Specifies that packet forwarding is to remain uninterrupted when OSPF services are restarted.
Enable Default Originate	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
Area ID	Type an area ID in the form of an IP address or decimal number.
Area Type	Select Normal or NSSA . 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.
Area Authentication	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"> ■ None No authentication is required. ■ Password With this choice, the password you specify in the Area Authentication Value field is included in the transmitted packet. ■ MD5 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 Area Authentication Value field.

- 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Enable Graceful Restart	Specifies that packet forwarding is to remain uninterrupted when BGP services are restarted.
Enable Default Originate	Allows the edge gateway to advertise itself as a default gateway to its BGP neighbors.
Local AS	<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
IP Address	Type the IP address of a BGP neighbor for this edge gateway.
Remote AS	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.
Weight	The default weight for the neighbor connection. Adjust as appropriate for your organization's needs.
Keep Alive Time	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
Hold Down Time	<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>
Password	<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>
BGP Filters	<p>Use this table to specify route filtering using a prefix list from this BGP neighbor.</p> <p>Caution 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 Keep to save each filter.</p> <ul style="list-style-type: none"> ■ Select the direction to indicate whether you are filtering traffic to or from the neighbor. ■ Select the action to indicate whether you are allowing or denying traffic. ■ Type the network that you want to filter to or from the neighbor. Type <code>ANY</code> or a network in a CIDR format. ■ Type the IP Prefix GE and IP Prefix LE to use the <code>le</code> and <code>ge</code> keywords in the IP prefix list.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Prefix Name	Select a specific IP prefix to apply this criteria to or select Any to apply the criteria to all network routes.
Learner Protocol	Select the protocol that is to learn routes from other protocols under this redistribution criteria.
Allow learning from	Select the types of networks from which routes can be learned for the protocol selected in the Learner Protocol list.
Action	Select whether to permit or deny redistribution from the selected types of networks.

- 6 Click **Save changes**.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 Acceleration Enabled 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>Note 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Type	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.
Enable SSL Passthrough	Click to enable SSL authentication to be passed through to the virtual server. Otherwise SSL authentication takes place at the destination address.
HTTP Redirect URL	(HTTP and HTTPS) Enter the URL to which traffic that arrives at the destination address should be redirected.
Persistence	<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"> ■ Source IP <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> ■ MSRDP <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> ■ SSL Session ID <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>
Cookie Name	(HTTP and HTTPS) If you specified Cookie 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"> Insert <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> Prefix <p>Select this option when your client does not support more than one cookie.</p> <p>Note 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> App Session 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.
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>Note 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 Insert X-Forwarded-For HTTP header for identifying the originating IP address of a client connecting to a Web server through the load balancer.</p> <p>Note Using this header is not supported if you enabled SSL passthrough.</p>
Enable Pool Side SSL	<p>(HTTPS Only) Select Enable Pool Side SSL 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>Note Select Enable Pool Side SSL 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. Note When set to Required , 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

- Open Edge Gateway Services.
 - From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - In the left panel, click **Edge Gateways**.
 - Click the radio button next to the name of the target edge gateway, and click **Services**.
- Navigate to **Load Balancer > Service Monitoring**.
- Click the **Create** () button.
- Enter a name for the service monitor.
- (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 Method .
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 Method 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 New Service Monitor dialog are activated or deactivated.

Option	Description
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).
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. Note When you select the POST method, you must specify a value for Send .
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. Note When Expected is not matched, the monitor does not try to match the Receive 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><extension>delay=2 critical=3 escape</extension></pre>

6 To preserve your changes, click **Keep**.

Example: Extensions Supported for Each Protocol

Table 7-1. 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. Note 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.

Table 7-1. Extensions for HTTP/HTTPS Protocols (continued)

Monitor Extension	Description
useragent= <i>STRING</i>	Sends the string in the HTTP header as <code>User Agent</code> .
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. <code>sticky</code> is like <code>follow</code> but stick to the specified IP address. <code>stickyport</code> ensures the port stays the same.
pagesize= <i>INTEGER:INTEGER</i>	Specifies the minimum and maximum page sizes required in bytes.
warning=DOUBLE	Specifies the response time in seconds to result in a warning status.
critical=DOUBLE	Specifies the response time in seconds to result in a critical status.

Table 7-2. 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 7-3. 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.

Table 7-3. Extensions for TCP Protocol (continued)

Monitor Extension	Description
delay= <i>INTEGER</i>	Waits the specified number of seconds between sending the string and polling for a response.
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

- 1 Open Edge Gateway Services.
 - a From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, and click **Services**.
- 2 Navigate to **Load Balancer > Pools**.
- 3 Click the **Create** () button.
- 4 Type a name and, optionally, a description for the load balancer pool.
- 5 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.

Option	Description
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.
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=<name></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=<url></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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Enable Virtual Server	Click to enable the virtual server.
Enable Acceleration	Click to enable acceleration.
Application Profile	Select an application profile to be associated with the virtual server.
Name	Type a name for the virtual server.
Description	Type an optional description for the virtual server.
IP Address	Type or browse to select the IP address that the load balancer listens on.
Protocol	Select the protocol that the virtual server accepts. You must select the same protocol used by the selected Application Profile .
Port	Type the port number that the load balancer listens on.
Default Pool	Choose the server pool that the load balancer will use.
Connection Limit	(Optional) Type the maximum concurrent connections that the virtual server can process.
Connection Rate Limit (CPS)	(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](#)

Secure Access Using Virtual Private Networks

You can configure the VPN capabilities that are provided by the NSX software for 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](#).

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
IP Range	Enter an IP address range for this IP pool, such as 127.0.0.1-127.0.0.9 . These IP addresses will be assigned to VPN clients when they authenticate and connect to the SSL VPN tunnel.
Netmask	Enter the netmask of the IP pool, such as 255.255.255.0 .
Gateway	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 IP Range field.
Description	(Optional) Enter a description for this IP pool.
Status	Select whether to activate or deactivate this IP pool.
Primary DNS	(Optional) Enter the name of the primary DNS server that will be used for name resolution for these virtual IP addresses.
Secondary DNS	(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
Network	Type the private network IP address in a CIDR format, such as 192169.1.0/24 .
Description	(Optional) Type a description for the network.
Send Traffic	<p>Specify how you want the VPN client to send the private network and Internet traffic.</p> <ul style="list-style-type: none"> ■ Over Tunnel <p>The VPN client sends the private network and Internet traffic over the SSL VPN-Plus activated edge gateway.</p> ■ Bypass Tunnel <p>The VPN client bypasses the edge gateway and sends the traffic directly to the private server.</p>
Enable TCP Optimization	<p>(Optional) To best optimize the Internet speed, when you select Over Tunnel for sending the traffic, you must also select Enable TCP Optimization. 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 Enable TCP Optimization eliminates the risk of this TCP-over-TCP problem occurring.</p> <p>Note When you activate TCP optimization:</p> <ul style="list-style-type: none"> ■ You must enter the port numbers for which to optimize the Internet traffic. ■ 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.
Ports	<p>When you select Over Tunnel, type a range of port numbers that you want opened for the remote user to access the internal servers, such as 20–21 for FTP traffic and 80–81 for HTTP traffic.</p> <p>To give unrestricted access to users, leave the field blank.</p>
Status	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
Enable password policy	Turn on enforcement of the password policy settings you configure here.
Password Length	Enter the minimum and maximum allowed number of characters for password length.
Minimum no. of alphabets	(Optional) Type the minimum number of alphabetic characters, that are required in the password.
Minimum no. of digits	(Optional) Type the minimum number of numeric characters, that are required in the password.
Minimum no. of special characters	(Optional) Type the minimum number of special characters, such as ampersand (&), hash tag (#), percent sign (%) and so on, that are required in the password.
Password should not contain user ID	(Optional) Enable to enforce that the password must not contain the user ID.
Password expires in	(Optional) Type the maximum number of days that a password can exist before the user must change it.
Expiry notification in	(Optional) Type the number of days prior to the Password expires in value at which the user is notified the password is about to expire.

- b (Optional) Enable and configure the account lockout policy.

Option	Description
Enable account lockout policy	Turn on enforcement of the account lockout policy settings you configure here.
Retry Count	Enter the number of times a user can try to access their account.
Retry Duration	Enter the time period in minutes in which the user account gets locked on unsuccessful login attempts. For example, if you specify the Retry Count as 5 and Retry Duration as 1 minute, the account of the user is locked after 5 unsuccessful login attempts within 1 minute.
Lockout Duration	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
Use this server for secondary authentication	(Optional) Specify whether to use the server as the second level of authentication.
Terminate session if authentication fails	(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
User ID	Enter the user ID.
Password	Enter a password for the user.
Retype Password	Reenter the password.
First name	(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
Profile Name	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.
Gateway	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 Add () button to add rows and type in their interface IP addresses or FQDNs, and ports.
Port	(Optional) To modify the port value from the displayed default, double-click the value and enter a new value.
Windows Linux Mac	Select the operating systems for which you want to create the installation packages.
Description	(Optional) Type a description for the user.
Enabled	Specify whether this package is activated or deactivated.

4 Select the installation parameters for Windows.

Option	Description
Start client on logon	Starts the SSL VPN client when the remote user logs in to their local system.
Allow remember password	Enables the client to remember the user password.
Enable silent mode installation	Hides installation commands from remote users.
Hide SSL client network adapter	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.
Hide client system tray icon	Hides the SSL VPN tray icon which indicates whether the VPN connection is active or not.
Create desktop icon	Creates an icon on the user desktop to invoke the SSL client.
Enable silent mode operation	Hides the window that indicates that installation is complete.
Server security certificate validation	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
Prevent multiple logon using same username	Turn on to restrict a remote user to having only one active login session under the same user name.
Compression	Turn on to enable TCP-based intelligent data compression and improve data transfer speed.
Enable Logging	Turn on to maintain a log of the traffic that passes through the SSL VPN gateway. Logging is enabled by default.
Force virtual keyboard	Turn on to require remote users to use a virtual (on-screen) keyboard only to enter login information.
Randomize keys of virtual keyboard	Turn on to have the virtual keyboard use a randomized key layout.
Session idle timeout	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.
User notification	Type the message to be displayed to remote users after they log in.
Enable public URL access	Turn on to allow remote users to access sites that are not explicitly configured by you for remote user access.
Enable forced timeout	Turn on to have the system disconnect remote users after the time period that you specify in the Forced timeout field is over.
Forced timeout	Type the timeout period in minutes. This field is displayed when Enable forced timeout 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 On the **IPsec VPN** tab, click **IPsec VPN Sites**.
- 2 Click the **Add** () button.

3 Configure the IPsec VPN connection settings.

Option	Action
Enabled	Enable this connection between the two VPN endpoints.
Enable perfect forward secrecy (PFS)	<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> <hr/> <p>Important 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>
Name	(Optional) Enter a name for the connection.
Local ID	<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>
Local Endpoint	<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>
Local Subnets	<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, 192.168.99.0/24.</p>
Peer ID	<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>
Peer Endpoint	<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> <hr/> <p>Note When NAT is configured for the peer, enter the public IP address that the device uses for NAT.</p>

Option	Action
Peer Subnets	<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, <code>192.168.99.0/24</code>.</p>
Encryption Algorithm	<p>Select the encryption algorithm type from the drop-down menu.</p> <p>Note The encryption type you select must match the encryption type configured on the remote site VPN device.</p>
Authentication	<p>Select an authentication. The options are:</p> <ul style="list-style-type: none"> ■ PSK <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> ■ Certificate <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 IPsec VPN tab's Global Configuration screen.</p>
Change Shared Key	<p>(Optional) When you are updating the settings of an existing connection, you can turn on this option on to make the Pre-Shared Key field available so that you can update the shared key.</p>
Pre-Shared Key	<p>If you selected PSK as the authentication type, type an alphanumeric secret string which can be a string with a maximum length of 128 bytes.</p> <p>Note 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	<p>(Optional) Enable this option to make the shared key visible in the screen.</p>
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>Note 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"> ■ <code>securelocaltrafficbyip=IPAddress</code> to redirect the edge gateway local traffic over the IPsec VPN tunnel. <p>This is the default value.</p> ■ <code>passthroughSubnets=PeerSubnet/IPAddress</code> to support overlapping subnets.

4 Click **Keep**.

5 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 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.

Procedure

1 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.

2 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.

3 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.

4 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.

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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
- b In the left panel, click **Edge Gateways**.
- c Click the radio button next to the name of the target edge gateway, 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.
Encryption Algorithm	Select the encryption algorithm for the communication between the server and the client.
Service Certificate Details	Click Change server certificate to select the certificate to be bound to the L2 VPN server. In the Change Server Certificate window, turn on Validate Server Certificate , select a server certificate from the list, and click OK .

- 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
Enabled	Enable this peer site.
Name	Enter a unique name for the peer site.
Description	(Optional) Type a description.

Option	Action
User ID	Enter the user name and password with which the peer site is to be authenticated.
Password	
Confirm Password	
Stretched Interfaces	User credentials on the peer site must be the same as the credentials on the client side.
	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.
Egress Optimization Gateway Address	(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
Server Address	Enter the IP address of the L2 VPN server to which this client is to be connected.
Server Port	Enter the L2 VPN server port to which the client should connect. The default port is 443.
Encryption Algorithm	Select the encryption algorithm for communicating with the server.
Stretched Interfaces	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.

Option	Description
Egress Optimization Gateway Address	(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.
User Details	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
Enable Secure Proxy	Select to enable the secure proxy.
Address	Enter the proxy server IP address.
Port	Enter the proxy server port.
User Name	Enter the proxy server authentication credentials.
Password	

- 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

The NSX 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 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Common Name	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.
Organization Unit	Use this field to differentiate between divisions within your VMware Cloud Director organization with which this certificate is associated. For example, Engineering or Sales.
Organization Name	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.
Locality	Enter the city or locality where your company is legally registered.
State or Province Name	Enter the full name (do not abbreviate) of the state, province, region, or territory where your company is legally registered.
Country Code	Enter the country name where your company is legally registered.
Private Key Algorithm	Enter 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. When FIPS mode is on, RSA key sizes must be greater or equal to 2048 bits. Note SSL VPN-Plus supports RSA certificates only.
Key Size	Enter the key size in bits. The minimum is 2048 bits.
Description	(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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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

The NSX 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. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Organization VDCs. c Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. d Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Edge Gateways. c Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. d Click the Grouping Objects tab.

- 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 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 an MAC set by using the **Grouping Objects** page. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ol style="list-style-type: none"> From the top navigation bar, under Resources, select Cloud Resources. In the left panel, click Organization VDCs. Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ol style="list-style-type: none"> From the top navigation bar, under Resources, select Cloud Resources. In the left panel, click Edge Gateways. Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. Click the Grouping Objects tab.

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 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 by using the **Grouping Objects** page. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Organization VDCs. c Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. d Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Edge Gateways. c Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. d Click the Grouping Objects tab.

- 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 by using the **Grouping Objects** page. To open this page, you must navigate either to the distributed firewall settings of the organization VDC, or to the services settings of an edge gateway that belongs to the organization VDC.

Procedure

- 1 Open the **Grouping Objects** page.

Option	Action
From the distributed firewall settings of the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Organization VDCs. c Select the radio button next to the name of the target organization virtual data center, and click Manage firewall. d Click the Grouping Objects tab.
From the services settings of an edge gateway on the organization VDC	<ol style="list-style-type: none"> a From the top navigation bar, under Resources, select Cloud Resources. b In the left panel, click Edge Gateways. c Select the radio button next to the name of an edge gateway that belongs to the target organization virtual data center, and click Services. d Click the Grouping Objects tab.

- 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.

View the Networks Use and IP Allocations on an Edge Gateway

You can view the networks on an edge gateway with information about their IP pool use and subnets. You can also view the IP address allocated to each network.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target edge gateway.
- 3 To view the external networks with information about their IP pool use and subnets, click the **External Networks > Networks & subnets** tab.
- 4 To view the external networks with information about their IP addresses and categories, click the **External Networks > IP allocations** tab.

Editing Edge Gateway Properties

Activate or Deactivate Distributed Routing on an Edge Gateway

After you activate VMware Cloud Director distributed routing on an edge gateway, the organization administrator can create many routed organization virtual data center networks with distributed interfaces connected to this edge gateway. Traffic on those networks is optimized for VM-to-VM communication.

Prerequisites

The backing NSX Manager instance is configured with an NSX Controller cluster. See the *NSX Administration Guide*.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Select the radio button next to the name of the target edge gateway, and click **Enable distributed routing** or **Disable distributed routing**.
- 4 To confirm, click **OK**.

Modify the External Networks and the Edge Gateway Settings

To modify the external networks and the edge gateway settings, you can use the **Edit edge gateway** wizard, which contains the same pages as the wizard that you used to create the edge gateway.

You can modify the settings that you configured when adding the edge gateway. See [Add an NSX Data Center for vSphere Edge Gateway](#) .

To modify the distributed routing setting, see [Activate or Deactivate Distributed Routing on an Edge Gateway](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the radio button next to the name of the edge gateway that you want to modify, and click **Edit**.
- 4 To modify the edge gateway settings, go through the pages of the **Edit edge gateway** wizard by clicking **Next**, and, on the **Ready to Complete** page, click **Finish**.

Edit the General Settings of an Edge Gateway

You can modify the name and the description of an edge gateway, activate or deactivate FIPS mode and high availability state, and change the edge gateway size configuration.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target edge gateway.
- 3 On the **General** tab, in the upper right corner, click **Edit**.
- 4 (Optional) Edit the name and the description of the edge gateway.
- 5 (Optional) Turn on or off each general edge gateway settings.

General Setting	Description
FIPS Mode	Configures the edge gateway to use NSX FIPS mode.
High Availability	Enables automatic failover to a backup edge gateway.

- 6 (Optional) Change the edge gateway configuration for your system resources.

Configuration	Description
Compact	Requires less memory and fewer compute resources.
Large	Provides increased capacity and performance than the Compact configuration. Large and X-Large configurations provide identical security functions.

Configuration	Description
X-Large	Used for environments that have a load balancer with large numbers of concurrent sessions.
Quad Large	Used for high throughput environments. Requires a high connection rate.

7 To confirm the changes, click **Save**.

Edit the Default Gateway of an Edge Gateway

You can change the network that an edge gateway uses as a default gateway.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target edge gateway.
- 3 On the **External Networks > Default gateway** tab, in the upper right corner, click **Edit**.
- 4 (Optional) Configure a network as the default gateway.
 - a Turn on the **Configure default gateway** toggle.
 - b Select the radio button next to the name of the target external network, and select the radio button next to the target IP address.
 - c (Optional) Turn on the **Use default gateway for DNS Relay** toggle.
- 5 To confirm the changes, click **Save**.

Edit the IP Settings of an Edge Gateway

You can modify the IP settings for external networks on an edge gateway.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target edge gateway.
- 3 On the **External Networks > IP settings** tab, click **Edit**.
- 4 For each network on the edge gateway, in the **IP Addresses** cell, enter an IP address or leave the cell blank.

If you do not enter an IP address for a network, the system assigns an arbitrary IP address to this network.
- 5 To confirm the changes, click **Save**.

Edit the Suballocated IP Pools on an Edge Gateway

You can suballocate multiple static IP pools from the available IP pools of an external network on an edge gateway.

Note Allocating IP addresses to an edge gateway through sub-allocation is a process where the provider assigns ownership of IP addresses to the gateway. VMware Cloud Director automatically configures the appropriate gateway interface with the secondary addresses during the sub-allocation process, which can cause IP address conflicts if any of the IP addresses are used outside of VMware Cloud Director.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target edge gateway.
- 3 Click the **External Networks > Sub-allocated IP pools** tab.

You can see the current suballocated IP pools for each external network on this edge gateway.

- 4 Click the radio button next to the name of an external network, and click **Edit**.

You can see the available IP pools for this external network, and the current suballocated IP pools if configured.

- 5 Edit the suballocated IP pools for this external network, and click **Save**.

You can add, modify, and remove IP addresses and ranges from the ranges of the available IP pools.

Results

The system combines overlapping IP ranges.

Edit the Rate Limits on an Edge Gateway

You can configure the inbound and outbound rate limits for each external network on the edge gateway.

Rate limits apply only to external networks backed by distributed port groups with static binding.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target edge gateway.
- 3 On the **External Networks > Rate limits** tab, in the upper right corner, click **Edit**.

You can see the current rate limits for each external network on this edge gateway.

- 4 Edit the rate limits, and click **Save**.

For each external network on the edge gateway, you can activate or deactivate the rate limits, and you can change the incoming and outgoing rates.

Redeploy an Edge Gateway

You can delete and deploy a new edge gateway appliance with the latest configurations.

If edge services are not working as expected, you can redeploy the edge gateway appliance.

When you redeploy an edge gateway, VMware Cloud Director deletes it and recreates it with the latest configurations.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the radio button next to the name of the target edge gateway, and click **Redeploy**.
- 4 To confirm, click **OK**.

Results

The edge gateway virtual machine is replaced with a new virtual machine and all services are restored.

Delete an Edge Gateway

You can remove an edge gateway from the organization virtual data center.

Prerequisites

Delete all organization virtual data center networks that use the target edge gateway.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the radio button next to the name of the target edge gateway, and click **Delete**.
- 4 To confirm, click **Delete**.

Statistics and Logs for an Edge Gateway

You can view statistics and logs for an edge gateway.

View Statistics

You can view statistics on the **Edge Gateway Services** screen.

Procedure

- 1 Open Edge Gateway Services.
 - a From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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
Connections	The Connections screen provides operational visibility. The screen displays graphs for the traffic flowing through the interfaces of the selected edge gateway and connection statistics for the firewall and load balancer services. Select the period for which you want to view the statistics.
IPsec VPN	The IPsec VPN screen displays the IPsec VPN status and statistics, and status and statistics for each tunnel.
L2 VPN	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

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

Procedure

- 1 Open Edge Gateway Services.
 - a From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, 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 Edge Gateway

You can enable SSH command-line access to an edge gateway.

Procedure

- 1 Open Edge Gateway Services.
 - a From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, and click **Services**.
- 2 Click the **Edge Settings** tab.

3 Configure the SSH settings.

Option	Description
Username	Enter the credentials for the SSH access to this edge gateway.
Password	By default, the SSH user name is admin .
Retype Password	
Password Expiry	Enter the expiration period for the password, in days.
Login Banner	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.

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.

This chapter includes the following topics:

- [Dedicated External Networks](#)
- [Add an NSX-T Data Center Edge Gateway](#)
- [Add an IP Set to an NSX-T Data Center Edge Gateway](#)
- [Add an NSX-T Data Center Edge Gateway Firewall Rule](#)
- [Add an SNAT or a DNAT Rule to an NSX-T Edge Gateway](#)
- [Configure a DNS Forwarder Service on an NSX-T Edge Gateway](#)
- [Edit the IP Allocations of an NSX-T Edge Gateway](#)
- [Quick IP Allocation](#)
- [Create Custom Application Port Profiles](#)
- [IPsec Policy-Based VPN for NSX-T Data Center Edge Gateways](#)
- [Configure Dedicated External Network Services](#)
- [Managing NSX Advanced Load Balancing on an NSX-T Data Center Edge Gateway](#)

Dedicated External Networks

To provide a fully routed network topology in a virtual data center, you 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 no other edge gateways can connect to the external network.

A tier-0 logical router or a VRF-lite 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.

Dedicating an external network to an edge gateway provides tenants with additional edge gateway services, such as route advertisement management and border gateway protocol (BGP) configuration.

The tenant can decide which of the tenant 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.

You can dedicate an external network to an NSX-T Data Center edge gateway either during the edge gateway creation or later, by editing the edge gateway general settings.

Add an NSX-T Data Center Edge Gateway

An NSX-T Data Center 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 firewall.

Prerequisites

For information about the system requirements for deploying an NSX-T Data Center edge gateway, see the *NSX-T Data Center Administration Guide*.

Starting with version 10.1, VMware Cloud Director supports a dedicated external network configuration. Dedicating an external network to an edge gateway provides tenants with additional edge gateway services, such as route advertisement management and border gateway protocol (BGP) configuration. For more information, see [Dedicated External Networks](#).

VMware Cloud Director supports basic NSX-T Data Center edge cluster configuration. For more information on NSX edge clusters, see *NSX-T Data Center Installation Guide*.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click **New**.
- 4 Select the NSX-T Data Center-backed organization VDC on which you want to create the edge gateway, and click **Next**.
- 5 Enter a name and, optionally, a description for the new edge gateway.
- 6 To enable BGP and route advertisement for the edge gateway, toggle on the **Dedicated External Network** option, and click **Next**.
- 7 Select an external network to which the new edge gateway connects and click **Next**.

If you toggled on the **Dedicated External Network** option, other edge gateways cannot access this external network.

- 8 Select an edge cluster on which to deploy the edge gateway and click **Next**.

If you want to run the edge gateway services on an edge cluster that is different from the one associated with the external network, you can configure the edge gateway to use a different edge cluster.

- Use the edge cluster of the external network to which the edge gateway is connected.
 - Select from a list of edge clusters available to the organization VDC on which you are deploying the edge gateway.
- 9 (Optional) Edit the IP addresses or IP address ranges that are allocated to the edge gateway, and click **Next**.

- 10 Review the **Ready to Complete** page, and click **Finish**.

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the NSX-T edge gateway.
- 4 Under **Security**, click **IP Sets** tab and click **New**.
- 5 Enter a name and, optionally, a description for the IP set.
- 6 Enter an IP address or an IP addresses range for the virtual machines that the IP set includes, and click **Add**.
- 7 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the edge gateway.
- 4 If the **Firewall** screen is not already visible under the Services section, click the **Firewall** tab.
- 5 Click **Edit Rules**.
- 6 Click the **New On Top** button.

A row for the new rule is added above the selected rule.

- 7 Configure the firewall rule.

Option	Description
Name	Enter a name for the rule.
State	To enable the rule upon creation, turn on the State toggle.
Applications	(Optional) To select a specific port profile to which the rule applies, turn on the Applications toggle and click Save .
Source	Select an option and click Keep . <ul style="list-style-type: none"> ■ To allow or deny traffic from any source address, toggle on Any Source. ■ To allow or deny traffic from specific firewall groups, select the firewall groups from the list.
Destination	Select an option and click Keep . <ul style="list-style-type: none"> ■ To allow or deny traffic to any destination address, toggle on Any Destination. ■ To allow or deny traffic to specific firewall groups, select the firewall groups from the list.
Action	From the Action drop-down menu, select an option. <ul style="list-style-type: none"> ■ To allow traffic from or to the specified sources, destinations, and services, select Accept. ■ To block traffic from or to the specified sources, destinations, and services, without notifying the blocked client select Drop. ■ To block traffic from or to the specified sources, destinations, and services, and to notify the blocked client that traffic was rejected, select Reject.
IP Protocol	Select whether to apply the rule to IPv4 or IPv6 traffic.
Direction	Select the traffic direction to which to apply the rule. <p>Note In VMware Cloud Director 10.2.1 and later versions, this option is no longer available.</p>
Enable logging.	To have the address translation performed by this rule logged, turn on the Enable logging toggle.

- 8 Click **Save**.
- 9 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the edge gateway and, under **Services**, click **NAT**.
- 4 To add a rule, click **New**.

5 Configure an SNAT or NO SNAT rule (inside going outside).

Option	Description
Name	Enter a meaningful name for the rule.
Description	(Optional) Enter a description for the rule.
Interface type	From the drop-down menu, select SNAT or NO SNAT.
External IP	Depending on the type of rule that you are creating, choose one of the options. <ul style="list-style-type: none"> ■ If you are creating a SNAT rule, enter the public IP address of the edge gateway for which you are configuring the SNAT rule. ■ If you are creating a NO SNAT rule, leave the text box empty.
Internal IP	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.
Destination IP	(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.
Advanced Settings (Optional)	Click the Advanced Settings tab for some additional settings. <p>State</p> <p>To activate the rule upon creation, toggle on the State option.</p> <p>Logging</p> <p>To have the address translation performed by this rule logged, toggle on the Logging option.</p> <p>Priority</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>Firewall Match</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"> ■ To apply firewall rules to the internal address of a NAT rule, select Match Internal Address. ■ To apply firewall rules to the external address of a NAT rule, select Match External Address. ■ To skip applying firewall rules, select Bypass.

6 Configure a DNAT or NO DNAT rule (outside going inside).

Option	Description
Name	Enter a meaningful name for the rule.
Description	(Optional) Enter a description for the rule.

Option	Description
Interface type	From the drop-down menu, select DNAT or NO DNAT.
External IP	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.
External Port	(Optional) Enter a port into which the DNAT rule is translating for the packets inbound to the virtual machines.
Internal IP	Depending on the type of rule that you are creating, choose one of the options. <ul style="list-style-type: none"> ■ 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. ■ If you are creating a NO DNAT rule, leave the text box empty.
Application	(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.
Advanced Settings (Optional)	Click the Advanced Settings tab for some additional settings. <p>State</p> <p>To activate the rule upon creation, toggle on the State option.</p> <p>Logging</p> <p>To have the address translation performed by this rule logged, toggle on the Logging option.</p> <p>Priority</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>Firewall Match</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"> ■ To apply firewall rules to the internal address of a NAT rule, select Match Internal Address. ■ To apply firewall rules to the external address of a NAT rule, select Match External Address. ■ To skip applying firewall rules, select Bypass.

7 Click **Save**.

8 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the edge gateway and, under **IP Management**, click **DNS**.
- 4 In the **DNS Forwarder** section, click **Edit**.
- 5 To enable the DNS Forwarder service, turn on the **State** toggle.
- 6 Enter a name and, optionally, a description for the default DNS zone.
- 7 Enter one or more upstream server IP addresses, separated by a comma.
- 8 Click **Save**.
- 9 (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**.

Edit the IP Allocations of an NSX-T Edge Gateway

You can allocate multiple IP addresses of an external network to an edge gateway.

Procedure

- 1 Open Edge Gateway Services.
 - a From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, and click **Services**.
- 2 Click the edge gateway, and click **IP Allocations**.

In the IP management grids, you can see the IP addresses that are allocated to the edge gateway and the IP addresses that are currently in use by the edge gateway.

- 3 In the **Allocated IPs** section, click **IP Management**.

In the **IP Management** grid, you can view the IP usage for each of the external networks that are available for use by the edge gateway.

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

Results

The IP addresses are allocated to the edge gateway.

What to do next

View the IP addresses that are allocated to the edge gateway, add more IP addresses, or remove them as needed.

Quick IP Allocation

You can allocate IP addresses from an external network subnet to an edge gateway without entering specific IP addresses or IP address ranges by using quick IP allocation.

Procedure

- 1 Open Edge Gateway Services.
 - a From the top navigation bar, select **Resources**, and click the **Cloud Resources** tab.
 - b In the left panel, click **Edge Gateways**.
 - c Click the radio button next to the name of the target edge gateway, and click **Services**.

- 2 Click the edge gateway and click **IP Allocations**.

In the IP management grids, you can see the IP addresses that are allocated to the edge gateway and the IP addresses that are currently in use by the edge gateway.

- 3 In the **Allocated IPs** section, click **Quick IP Allocation**.
- 4 From the drop-down menu, select a subnet from which to assign IP addresses.

If multiple subnets are available, selecting **Any** results in the allocation of IP addresses from one or more subnets.

- 5 Enter the number of IP addresses to allocate to the edge gateway, and click **Save**.

The number must be less than the number of available IP addresses in the subnet that you selected.

Results

The IP addresses are allocated to the edge gateway.

What to do next

View the IP addresses that are allocated to the edge gateway, add more IP addresses, or remove them as needed.

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the edge gateway.
- 4 Under **Security**, click **Application Port Profiles**.
- 5 In the **Custom Applications** section, click **New**.
- 6 Enter a name and, optionally, a description for the application port profile.
- 7 Select a protocol from the drop-down menu.
- 8 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target 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 \checkmark .

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target 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
IKEv1	When you select this option, IPsec VPN initiates and responds to IKEv1 protocol only.
IKEv2	The default option. When you select this version, IPsec VPN initiates and responds to IKEv2 protocol only.
IKE-Flex	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
Copy	Copies the defragmentation bit from the inner IP packet to the outer packet.
Clear	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.

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 you have dedicated an external network to an NSX-T Data Center edge gateway in the organization. See [Dedicated External Networks](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target 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. You can edit these settings on the parent tier-0 gateway in NSX-T Data Center.

Prerequisites

- Verify that you have dedicated an external network to an NSX-T Data Center edge gateway in the organization. See [Dedicated External Networks](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target 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
Helper and graceful restart	<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>Note The edge gateway configuration applies to all BGP neighbors unless the neighbor-specific configuration overrides it.</p>
Helper only	<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>
Disable	<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 you have dedicated an external network to an NSX-T Data Center edge gateway in the organization. See [Dedicated External Networks](#).
- [Configure BGP General Settings](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target 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 you have dedicated an external network to an NSX-T Data Center edge gateway in the organization. See [Dedicated External Networks](#).
- 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**, and click the name of the target 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
Disable	Overrides the global edge gateway settings and deactivates graceful restart mode for this neighbor.
Helper only	Overrides the global edge gateway settings and configures graceful restart mode as Helper only for this neighbor.
Graceful restart and Helper	Overrides the global edge gateway settings and configures graceful restart mode as Graceful restart and Helper 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.

Managing NSX Advanced Load Balancing on an NSX-T Data Center Edge Gateway

As a **system administrator**, you enable load balancing on an NSX-T Data Center gateway and assign a service engine group to the edge gateway.

An **organization administrator** creates load balancer server pools and virtual services.

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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the NSX-T Data Center edge gateway on which you want to enable load balancing.
- 4 Under Load Balancer, click **General Settings**.
- 5 Click **Edit** and toggle on the **Load Balancer State** option.
- 6 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.
- 7 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the NSX-T Data Center edge gateway to which you want to assign a service engine group.
- 4 Under Load Balancer, click **Service Engine Groups**.
- 5 Click **Add**.
- 6 Select an available service engine group from the list.
- 7 Enter a number for the maximum number of virtual services that can be placed on the edge gateway.
- 8 Enter a number for the guaranteed virtual services available to the edge gateway.
- 9 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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the NSX-T Data Center edge gateway to which the service engine group is assigned.
- 4 Under Load Balancer, click **Service Engine Groups**.
- 5 Click **Edit**.
- 6 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.
- 7 Edit the number for the guaranteed virtual services available to the edge gateway.
- 8 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

- [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 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the NSX-T Data Center edge gateway for which you want to configure a load balancer pool.
- 4 Under Load Balancer, click **Pools**, and then click **Add**.

- 5 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
Least Connections	New connections are sent to the server that currently has the fewest connections.
Round Robin	New connections are sent to the next eligible server in the pool in a sequential order.
Fastest Response	New connections are sent to the server that provides the fastest response to new connections or requests.
Consistent Hash	New connections are distributed across the servers by using the IP address of the client to generate an IP hash key.
Least Load	New connections are sent to the server with the lightest load, regardless of the number of connections that server has.
Fewest Servers	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.
Random	The load balancer picks servers at random.
Fewest Tasks	Load is adaptively balanced, based on the server feedback.
Core Affinity	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.

6 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.

7 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.

8 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.

9 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

- [Enable Load Balancer on an NSX-T Data Center Edge Gateway](#).
- [Assign a Service Engine Group to an NSX-T Data Center Edge Gateway](#).
- [Add a Load Balancer Server Pool](#).

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, click **Edge Gateways**.
- 3 Click the NSX-T Data Center edge gateway on which you want to create a virtual service.
- 4 Under Load Balancer, click **Virtual Services**, and then click **Add**.
- 5 Enter a meaningful name and, optionally, a description, for the virtual service.
- 6 To activate the virtual service upon creation, toggle on the **Enabled** option.
- 7 Select a service engine group for the virtual service.
- 8 Select a load balancer pool for the virtual service.
- 9 Enter an IP address for the virtual service.
- 10 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.
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.

11 Click **Save**.

Managing Dedicated vCenter Server Instances

9

With dedicated vCenter Server instances, you can use VMware Cloud Director as a central point of management (CPOM) for your vSphere environments.

When you add a vCenter Server instance to VMware Cloud Director, you can specify the purpose of the instance.

Dedicated vCenter Server

The infrastructure of an attached vCenter Server instance is encapsulated as a Software-Defined Data Center (SDDC) and is fully dedicated to a single tenant. You create a dedicated vCenter Server instance by activating the tenant access for that instance. After you activate the tenant access, you can publish a dedicated vCenter Server instance to a tenant.

Shared vCenter Server

The provider can use different resource pools of the vCenter Server instance across multiple provider VDCs and then allocate those resource pools to different tenants. A shared vCenter Server instance cannot be published to tenants.

None

The vCenter Server instance does not have any specific purpose.

VMware Cloud Director can act as an HTTP proxy server for the dedicated vCenter Server instances and the vCenter Server instances that do not have a set purpose.

With dedicated vCenter Server instances, you can use VMware Cloud Director as a central point of management for all your vSphere environments.

- You can dedicate the resources of a vCenter Server instance to a single tenant by publishing the corresponding dedicated vCenter Server only to its organization. The tenant does not share these resources with other tenants. The tenant can access this dedicated vCenter Server instance by using a UI or API proxy without a VPN required.
- You can use VMware Cloud Director as a lightweight directory to register all your vCenter Server instances.
- You can use VMware Cloud Director as an API endpoint for all your vCenter Server instances.

You can activate the tenant access and mark a vCenter Server instance as dedicated, during or after the attachment of the target vCenter Server instance to VMware Cloud Director. See [Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance](#).

With an attached vCenter Server instance, you can create either a shared vCenter Server or a dedicated vCenter Server. If you created a shared vCenter Server instance, you cannot use this vCenter Server instance to create a dedicated vCenter Server, and the reverse.

You can create endpoints that tenants can use to access the underlying vSphere environment. By using their VMware Cloud Director accounts, users can log in to the UI or API of components with or without proxies.

Dedicated vCenter Server instances in VMware Cloud Director remove the requirement for vCenter Server to be publicly accessible. To control the access, you can activate and deactivate the tenant access to an SDDC in VMware Cloud Director.

An endpoint is the access point to a component from an SDDC, for example, a vCenter Server instance, an ESXi host, or an NSX Manager instance. You can connect an endpoint to a proxy. By activating and deactivating a proxy, you can allow and stop the tenant access through that proxy.

Starting with VMware Cloud Director 10.2, if you use the API to query the dedicated vCenter Server and proxy entities and your tenant configuration supports multisite associations, VMware Cloud Director returns a multisite response. The results are from all available associations.

Creating and Managing Dedicated vCenter Server Instances

To create and manage dedicated vCenter Server instances and proxies, you can use the Service Provider Admin Portal or the VMware Cloud Director OpenAPI. For VMware Cloud Director OpenAPI, see *Getting Started with VMware Cloud Director OpenAPI* at <https://code.vmware.com>.

Important VMware Cloud Director requires a direct network connection to each dedicated vCenter Server instance. If the vCenter Server instance uses an external Platform Services Controller, VMware Cloud Director requires a direct network connection to the Platform Services Controller as well.

To use VMware OVF Tool in a proxied dedicated vCenter Server, VMware Cloud Director requires a direct connection to each ESXi host.

1 Create a dedicated vCenter Server instance.

When you add a vCenter Server instance to the VMware Cloud Director environment, you can create a dedicated vCenter Server instance by activating the tenant access in the **Add vCenter Server** wizard. See [Add the vCenter Server Instance](#).

Creating a dedicated vCenter Server instance also creates a default endpoint for it. While attaching the vCenter Server instance, you can also create a proxy. However, the default endpoint is not connected to any proxy by default. You must edit the default endpoint or create a new one to connect it to a proxy. See [Create an Endpoint](#).

You can activate the tenant access of vCenter Server instances that are already added to VMware Cloud Director and do not have a specified use. See [Enable the Tenant Access of an Attached vCenter Server](#). Activating the tenant access makes the vCenter Server instance available to be published to tenants.

2 Add a proxy.

You can create a proxy either when you attach a vCenter Server instance to VMware Cloud Director or later. If the vCenter Server instance uses an external Platform Services Controller, VMware Cloud Director creates a proxy for the Platform Services Controller as well. With parent and child proxies, you can hide certain proxies from the tenants or you can activate and deactivate groups of child proxies through their parent proxies. For information on creating a proxy after you add a vCenter Server instance to VMware Cloud Director, see [Add a Proxy for Accessing the Underlying vCenter Server Resources](#).

You can edit, activate, deactivate, and delete proxies from the **Proxies** tab under **vSphere Resources**.

Note When you add a proxy to a dedicated vCenter Server instance, you must upload the certificate and the thumbprint, so that tenants can retrieve the certificate and the thumbprint if the proxied component uses self-signed certificates.

To view and manage certificates and certificate revocation lists (CRLs), see [Manage the Proxy Certificates and CRLs](#).

3 Get the certificate and the thumbprint of the created proxies, and verify that the certificate and the thumbprint are present and correct. See [Manage the Proxy Certificates and CRLs](#).

4 Publish the dedicated vCenter Server instance to one or more organizations.

You can publish a dedicated vCenter Server instance to a tenant and make it visible in the VMware Cloud Director Tenant Portal. In most cases, one vCenter Server instance should be published only to one tenant. See [Publish a Dedicated vCenter Server](#).

5 To enable the tenants to access the dedicated vCenter Server instances and proxies from the VMware Cloud Director Tenant Portal, you must publish the **CPOM extension** plug-in to their organizations. See [Publish or Unpublish a Plug-in from an Organization](#).

This chapter includes the following topics:

- [Enable the Tenant Access of an Attached vCenter Server](#)
- [Publish a Dedicated vCenter Server](#)

Enable the Tenant Access of an Attached vCenter Server

You can enable the tenant access of vCenter Server instances that are already added to VMware Cloud Director and do not have a specified use. Enabling the tenant access creates a dedicated vCenter Server instance and makes it available to be published to tenants.

With an attached vCenter Server instance, you can create either a shared vCenter Server or a dedicated vCenter Server. If you created a shared vCenter Server instance, and you want to use it as a dedicated vCenter Server, you must first delete all provider virtual data centers (VDCs) that are using the resources of the vCenter Server instance. Deleting all provider VDCs linked to the shared vCenter Server instance changes its status to None.

Prerequisites

Verify that you have in your environment at least one attached vCenter Server that is not dedicated or shared.

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.
- 3 Select a vCenter Server without a specified purpose in the **Usage** column.
- 4 Click **Enable Tenant Access**.

What to do next

[Publish a Dedicated vCenter Server](#).

Publish a Dedicated vCenter Server

You can publish a dedicated vCenter Server to a tenant and make it visible through the VMware Cloud Director Tenant Portal. By default, one vCenter Server should be published only to one tenant.

By default, an SDDC is a vCenter Server instance that you dedicate to a single tenant by publishing the corresponding dedicated vCenter Server instance only to its organization. The tenant does not share the dedicated vCenter Server instance resources with other tenants. Publishing a dedicated vCenter Server instance to multiple tenants violates the tenancy boundaries. However, sometimes a tenant must have access to multiple dedicated vCenter Server instances. In these cases, you can publish a dedicated vCenter Server instance to multiple tenants.

Prerequisites

- Verify that you have at least one vCenter Server instance with enabled tenant access in your VMware Cloud Director environment. See [Chapter 9 Managing Dedicated vCenter Server Instances](#).

Procedure

- 1 From the top navigation bar, under **Resources**, click **Infrastructure Resources**.
- 2 In the left panel, select **vCenter Server Instances**.

- 3 Select a vCenter Server with enabled tenant access.

The vCenter Server instances with enabled tenant access have a Dedicated value in the **Usage** column.

- 4 Click **Manage Tenants**.

- 5 Select the tenant or tenants to which you want to publish the vCenter Server instance.

Deselecting a tenant from the list, unpublishes the vCenter Server.

- 6 Click **Save**.

What to do next

To enable users to access the dedicated vCenter Server instances and the proxies from the VMware Cloud Director Tenant Portal, you must publish the **CPOM extension** plug-in to their organizations. See [Publish or Unpublish a Plug-in from an Organization](#).

Managing System Administrators and Roles

10

By using the VMware Cloud Director Service Provider Admin Portal, you can add system 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.

Note Starting with VMware Cloud Director 9.5, service providers can create provider roles and manage provider users and groups by using the VMware Cloud Director Service Provider Admin Portal or by using the vCloud OpenAPI. For information about managing provider roles, users, and groups, see the *VMware Cloud Director Service Provider Admin Portal Guide*. To examine the vCloud OpenAPI documentation, go to https://vCloud_Director_IP_address_or_host_name/docs.

This chapter includes the following topics:

- [Managing Rights and Roles](#)
- [Managing Provider Users and Groups](#)

Managing Rights and Roles

A right is the fundamental unit of access control in VMware Cloud Director. A role associates a role name with a set of rights. Each organization can have different rights and roles.

VMware Cloud Director uses roles and their associated rights to determine whether a user or group is authorized to perform an operation. Many of the procedures documented in the VMware Cloud Director guides include a prerequisite role. These prerequisites assume that the named role is the unmodified predefined role or a role that includes an equivalent set of rights.

System administrators can use rights bundles and global tenant roles to manage the rights and roles that are available to each organization.

After you install VMware Cloud Director, the system contains only the System Rights Bundle, which includes all rights that are available in the system. The System Rights Bundle is not published to any organization. The system also contains built-in global tenant roles that are published to all organizations. For information about the predefined roles, see [Predefined Roles and Their Rights](#).

After you upgrade VMware Cloud Director from version 9.1 or earlier, in addition to the System Rights Bundle, the system contains a Legacy Rights Bundle for each existing organization. Each Legacy Rights Bundle includes the rights that are available in the associated organization at the time of the upgrade and is published only to this organization.

Note To begin using the rights bundles model for an existing organization, you must delete the corresponding Legacy Rights Bundle.

If you upgraded VMware Cloud Director from version 9.1 or earlier, the existing role templates are published to all organizations as global tenant roles, and the existing roles that are unlinked from role templates are available as tenant-specific roles to their organizations.

Rights Terminology

Right

Each right provides view or manage access to a particular object type in VMware Cloud Director. Rights belong to different categories depending on the objects to which they relate, for example, vApp, Catalog, Organization, and so on. The Provider organization contains all rights available in the system. The system administrator defines the rights that are available to each organization. You cannot create or modify the rights included in VMware Cloud Director.

Rights Bundle

System administrators can use rights bundles to manage the rights that are available to each organization. A rights bundle is a set of rights that the system administrator can publish to one or more organizations. The system administrator can create and publish rights bundles that correspond to tiers of service, separately monetizable functionality, or any other arbitrary rights grouping. Only system administrators can view and manage the rights bundles. You can publish multiple bundles to the same organization.

Organization Rights

Organization rights are the full set of rights that are available to an organization. Organization rights can comprise multiple rights bundles, but the organization administrators and users see a flat set of rights that they can use to create and modify tenant-specific roles.

Roles Terminology

Role

A role is a set of rights that is assignable to one or more users and groups. When you create or import a user or group, you must assign it a role.

Provider Roles

Provider roles are the set of roles that are available only to the Provider organization. Provider roles can be assigned only to Provider users. System administrators can create custom provider roles.

Tenant Roles

Tenant roles are the set of roles available to an organization.

System administrators can create and edit global tenant roles and publish them to one or more organizations. Global tenant roles can be assigned to tenant users in the organizations to which they are published. Organization administrators cannot edit global tenant roles.

Note Tenant users can use only those rights from their roles that are published to their organizations.

Tenant-Specific Roles

Organization administrators can create and edit tenant-specific roles, which are local to their organizations. Tenant-specific roles can be assigned only to tenant users in the organization to which they belong. Tenant-specific roles can contain a subset of the organization rights only.

For information about managing tenant-specific roles, see *VMware Cloud Director Tenant Portal Guide*.

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 [System Administrator Rights](#). 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

A **System Administrator** can use the Service Provider Admin Portal to view the list of rights included in a role.

- 1 In the top navigation bar, click **Administration**.
- 2 From the left panel under **Provider Access Control**, select **Roles**.
- 3 Click the name of the role you want to view.

An **Organization Administrator** can use the Service Provider Admin Portal or the Cloud Director OpenAPI to view the rights in a role or create roles local to the organization.

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](#).

System Administrator Rights

The **System Administrator** role exists only in the provider organization. By default the **System Administrator** role has all VMware Cloud Director rights.

The **System Administrator** role has all VMware Cloud Director rights. This list consists of the rights available only to **System Administrators**. The **System Administrator** role has also the [Rights in Predefined Global Tenant Roles](#).

Table 10-1. Rights Available by Default Only to System Administrators

New in this release	Right Name
	Access All Organization VDCs
	Access Control List: Manage
	Access Control List: View
	Additional Services: Execute Workflows
	Additional Services: View Running Workflows

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Additional Services: View Workflows
	Adopt Resource Pool: View
✓	Advisory Definitions: Create and Delete
✓	Advisory Definitions: Read
	Alternate Admin Entity: View
	AMQP Settings: Manage
	AMQP Settings: View
	API Explorer: View
	Catalog: Add vApp from My Cloud
	Catalog: Change Owner
	Catalog: Create / Delete a Catalog
	Catalog: Edit Properties
	Catalog: Import Media from vSphere
	Catalog: Publish
	Catalog: Shadow VM View
	Catalog: Sharing
	Catalog: VCSP Publish Subscribe
	Catalog: VCSP Publish Subscribe Caching
	Catalog: View ACL
	Catalog: View Private and Shared Catalogs
	Catalog: View Published Catalogs
	Cell Configuration: View
	Certificate Library: Manage
	Certificate Library: View
	Cloud Tunnel Server: Manage
	Cloud Tunnel Server: View
	Content Library System Settings: Manage
	Content Library System Settings: View

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Custom entity: Create custom entity definitions
	Custom entity: Delete custom entity definitions
	Custom entity: Edit custom entity definitions
	Custom entity: View all custom entity instances in org
	Custom entity: View custom entity definitions
	Custom entity: View custom entity instance
	Datastore: Delete
	Datastore: Edit
	Datastore: Enable or Disable
	Datastore: Open in vSphere
	Datastore: View
	Direct Org vDC Network: Manage
	Distributed Virtual Switch: Open in vSphere
	Edge Cluster: Manage
	Edge Cluster: View
	Extension Service API Definition: Manage
	Extension Service API Definition: View
	Extension Services: View
	Extensions: View
	External Service: Manage
	External Service: View
✓	General ACL: Manage
✓	General ACL: View
	General: Administrator Control
	General: Administrator View
	General: Send Notification
	General: View Error Details
	Global Role: Edit

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Global Role: View
	Group / User: View
	Host: Enable or Disable
	Host: Manage
	Host: Open in vSphere
	Host: Prepare or Unprepare
	Host: Repair
	Host: Upgrade
	Host: 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
	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
	Kerberos Settings: Manage
	Kerberos Settings: View
	LDAP Settings: Manage
	LDAP Settings: View
	License Report: View
✓	Load Balancer Controller: Edit
✓	Load Balancer Controller: View
✓	Load Balancer Service Engine Group Assignment: Edit
✓	Load Balancer Service Engine Group Assignment: View

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
✓	Load Balancer Service Engine Group: Edit
✓	Load Balancer Service Engine Group: View
	Localization Resources: Manage
	Network Pool: Create or Delete
	Network Pool: Edit
	Network Pool: Open in vSphere
	Network Pool: Repair
	Network Pool: View
	NSX-T: Edit
	NSX-T: View
	Object Extensions: Manage
	Object Extensions: View
	Organization Network: Create or Delete
	Organization Network: Edit Properties
	Organization Network: Open in vSphere
	Organization Network: View
✓	Organization Quotas: Manage
	Organization vDC Compute Policy: Admin View
	Organization vDC Compute Policy: Manage
	Organization vDC Compute Policy: View
	Organization vDC Distributed Firewall: Configure Rules
	Organization vDC Distributed Firewall: Enable/Disable
	Organization vDC Distributed Firewall: View Rules
	Organization vDC Gateway: Configure BGP Routing
	Organization vDC Gateway: Configure DHCP
	Organization vDC Gateway: Configure DNS
	Organization vDC Gateway: Configure ECMP Routing
	Organization vDC Gateway: Configure Firewall

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Organization vDC Gateway: Configure IPsec VPN
	Organization vDC Gateway: Configure L2 VPN
	Organization vDC Gateway: Configure Load Balancer
	Organization vDC Gateway: Configure NAT
	Organization vDC Gateway: Configure OSPF Routing
	Organization vDC Gateway: Configure Remote Access
	Organization vDC Gateway: Configure Route Advertisement
✓	Organization vDC Gateway: Configure SLAAC Profile
	Organization vDC Gateway: Configure SSL VPN
	Organization vDC Gateway: Configure Static Routing
	Organization vDC Gateway: Configure Syslog
	Organization vDC Gateway: Configure System Logging
	Organization vDC Gateway: Convert to Advanced Networking
	Organization vDC Gateway: Create
	Organization vDC Gateway: Delete
	Organization vDC Gateway: Distributed Routing
	Organization vDC Gateway: Import
	Organization vDC Gateway: Modify Form Factor
	Organization vDC Gateway: Update
	Organization vDC Gateway: Update Properties
	Organization vDC Gateway: Upgrade
	Organization vDC Gateway: View
	Organization vDC Gateway: View BGP Routing
	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 L2 VPN

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Organization vDC Gateway: View Load Balancer
	Organization vDC Gateway: View NAT
	Organization vDC Gateway: View OSPF Routing
	Organization vDC Gateway: View Remote Access
	Organization vDC Gateway: View Route Advertisement
✓	Organization vDC Gateway: View SLAAC Profile
	Organization vDC Gateway: View SSL VPN
	Organization vDC Gateway: View Static Routing
✓	Organization vDC Kubernetes Policy: Edit
	Organization vDC Named Disk: Change Owner
	Organization vDC Named Disk: Create
	Organization vDC Named Disk: Delete
	Organization vDC Named Disk: Edit Properties
	Organization vDC Named Disk: View Encryption Status
	Organization vDC Named Disk: View Properties
	Organization vDC Network: Edit Properties
	Organization vDC Network: Import
	Organization vDC Network: View
	Organization vDC Resource Pool: Open in vSphere
	Organization vDC Resource Pool: View
✓	Organization vDC Shared Named Disk: Create
	Organization vDC Storage Policy: Edit
	Organization vDC Storage Policy: Enable or Disable
	Organization vDC Storage Policy: Open in vSphere
	Organization vDC Storage Policy: Remove
	Organization vDC Storage Policy: View Capabilities
	Organization vDC Storage Profile: Set Default
	Organization vDC: Create

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Organization vDC: Delete
	Organization vDC: Edit ACL
	Organization vDC: Enable or Disable
	Organization vDC: Extended Edit
	Organization vDC: Extended View
	Organization vDC: Manage Firewall
	Organization vDC: Simple Edit
	Organization vDC: User View
	Organization vDC: View ACL
	Organization VDC: view metrics
	Organization vDC: VM-VM Affinity Edit
	Organization: Activate or Deactivate
	Organization: Create or Delete
	Organization: Edit Association Settings
	Organization: Edit Federation Settings
	Organization: Edit LDAP Settings
	Organization: Edit Leases Policy
	Organization: Edit Limits
	Organization: Edit Name
	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: Migrate Tenant Storage
	Organization: Perform Administrator Queries
	Organization: Use Provider LDAP as Tenant

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Organization: View
	Organization: view metrics
	Port Group: Open in vSphere
	Preference: Manage preference definition
	Provider Network: Create or Delete
	Provider Network: Edit
	Provider Network: Open in vSphere
	Provider Network: View
	Provider vDC Compute Policy: Manage
	Provider vDC Compute Policy: View
	Provider vDC Resource Pool: Migrate VMs
	Provider vDC Resource Pool: Open in vSphere
	Provider vDC Resource Pool: View
	Provider vDC Storage Policy: Edit
	Provider vDC Storage Policy: Enable or Disable
	Provider vDC Storage Policy: Open in vSphere
	Provider vDC Storage Policy: Remove
	Provider vDC Storage Policy: View
	Provider vDC: Add Resource Pool
	Provider vDC: Create or Delete
	Provider vDC: Delete Resource Pool
	Provider vDC: Edit
	Provider vDC: Enable or Disable
	Provider vDC: Enable or Disable Resource Pool
	Provider vDC: Enable vSphere VXLAN
	Provider vDC: Merge
	Provider vDC: View
✓	Quota Policy Capabilities: View

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
✓	Quota Policy: Manage
✓	Quota Policy: View
	Reload VM: Manage
	Resource Class Action: Manage
	Resource Class Action: View
	Resource Pool: Open
	Resource Pool: Open in vSphere
	Resource Pool: View
	Right: Manage
	Right: View
	Rights Bundle: Edit
	Rights Bundle: View
	Role: Create, Edit, Delete, or Copy
	SDDC: Manage
	SDDC: Manage Proxy
	SDDC: View
	Selector Extensions: Manage
	Selector Extensions: View
	Service Apps: Manage
	Service Apps: View
	Service Authorization: Manage
	Service Configuration: Manage
	Service Configuration: View
	Service Library: Create service libraries
	Service Library: Delete services from the service library
	Service Library: Edit service metadata
	Service Library: Edit the contents of a service
	Service Library: View service libraries

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Service Link: Manage
	Service Link: View
	Service Resource Type: Manage
	Service Resource Type: View
	Service Resource: Manage
	Service Resource: View
	Shared Org vDC Network: Manage
	Site: Edit
	Site: View
	SSL Settings: View
✓ (Available in version 10.2.2 and later)	SSL Settings: Manage
✓	SSL: Test Connection
	Stranded Item: Manage
	Stranded Item: View
✓ (Available in version 10.2.2 and later)	Supported Storage Entity Type: Manage
	System Operations: Execute System Operations
	System Organization: Manage
	System Organization: View
	System Settings: Manage
	System Settings: View
✓	Tanzu Kubernetes Guest Cluster: Administrator Full Control
✓	Tanzu Kubernetes Guest Cluster: Administrator View
✓	Tanzu Kubernetes Guest Cluster: Edit
✓	Tanzu Kubernetes Guest Cluster: Full Control
✓	Tanzu Kubernetes Guest Cluster: View
	Task: Resume, Abort, or Fail

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	Task: Update
	Task: View Tasks
	Token: Manage
	Token: Manage All
	Truststore: Manage
	Truststore: View
	UI Plugins: Define, Upload, Modify, Delete, Associate or Disassociate
	UI Plugins: View
	UI Portal Branding: Manage
	vApp Template / Media: Copy
	vApp Template / Media: Create / Upload
	vApp Template / Media: Edit
	vApp Template / Media: View
	vApp Template: Add to My Cloud
	vApp Template: Change Owner
	vApp Template: Download
	vApp Template: Force storage lease expiration
	vApp Template: Import
	vApp Template: Open in vSphere
	vApp: Allow All Extra Config
	vApp: Allow Ethernet Coalescing Extra Config
	vApp: Allow Latency Extra Config
	vApp: Allow Matching Extra Config
	vApp: Allow NUMA Node Affinity Extra Config
	vApp: Change Owner
	vApp: Copy
	vApp: Create / Reconfigure
	vApp: Delete

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	vApp: Download
	vApp: Edit Properties
	vApp: Edit VM Compute Policy
	vApp: Edit VM CPU
	vApp: Edit VM CPU and Memory reservation settings in all VDC types
	vApp: Edit VM Hard Disk
	vApp: Edit VM Memory
	vApp: Edit VM Network
	vApp: Edit VM Properties
	vApp: Enter/Exit Maintenance Mode
	vApp: Force runtime lease expiration
	vApp: Force storage lease expiration
	vApp: Import Options
	vApp: Maintenance manage
	vApp: Manage VM Password Settings
	vApp: Open in vSphere
	vApp: Power Operations
	vApp: Shadow VM View
	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 Check Compliance
	vApp: VM Migrate, Force Undeploy, Relocate, Consolidate

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	VAPP_VM_METADATA_TO_VCENTER
	VCD Extension: Register, Unregister, Refresh, Associate or Disassociate
	VCD Extension: View
	vCenter: Attach or Detach
	vCenter: Enable or Disable
	vCenter: Open in vSphere
	vCenter: Refresh
	vCenter: View
	vDC Group: Configure
✓	vDC Group: Configure Logging
	vDC Group: View
	VDC Template: ACL manage
	VDC Template: Extended View
	VDC Template: Instantiate
	VDC Template: Manage
	VDC Template: View
	VMC: Register SDDC
✓	VMWARE:NATIVECLUSTER: Administrator Full Control
✓	VMWARE:NATIVECLUSTER: Administrator View
✓	VMWARE:NATIVECLUSTER: Edit
✓	VMWARE:NATIVECLUSTER: Full Control
✓	VMWARE:NATIVECLUSTER: View
	vRealize Orchestrator: Publish and Unpublish Workflows to Tenants
	vRealize Orchestrator: Register and Unregister vRealize Orchestrator Servers
	vRealize Orchestrator: View Registered vRealize Orchestrator Servers
	vSphere Server: Manage
	vSphere Server: Manage Proxy

Table 10-1. Rights Available by Default Only to System Administrators (continued)

New in this release	Right Name
	vSphere Server: Manage Proxy Configuration
	vSphere Server: View

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: Create / Delete a Catalog	✓	✓			
	Catalog: Edit Properties	✓	✓			
	Catalog: Publish	✓	✓			
	Catalog: Sharing	✓	✓			
	Catalog: VCSP Publish Subscribe	✓	✓			
	Catalog: View ACL	✓	✓			
	Catalog: View Private and Shared Catalogs	✓	✓	✓		
	Catalog: View Published Catalogs	✓				
	Certificate Library: Manage	✓				
	Certificate Library: View	✓				
	Custom entity: View all custom entity instances in org	✓				
	Custom entity: View custom entity instance	✓				
	General: Administrator Control	✓				
	General: Administrator View	✓				

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	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	✓				
	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	✓				

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	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	✓				
	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 Named Disk: Change Owner	✓	✓			
	Organization vDC Named Disk: Create	✓	✓	✓		
	Organization vDC Named Disk: Delete	✓	✓	✓		
	Organization vDC Named Disk: Edit Properties	✓	✓	✓		
	Organization vDC Named Disk: View Encryption Status	✓		✓		

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Organization vDC Named Disk: View Properties	✓	✓	✓	✓	
	Organization vDC Network: Edit Properties	✓				
	Organization vDC Network: View	✓		✓		
	Organization vDC Storage Policy: View Capabilities	✓				
	Organization vDC Storage Profile: Set Default	✓				
	Organization vDC: Edit ACL	✓				
	Organization vDC: Manage Firewall	✓				
	Organization vDC: Simple Edit	✓				
	Organization vDC: User View	✓	✓			
	Organization vDC: View ACL	✓				
	Organization VDC: view metrics	✓				
	Organization vDC: VM-VM Affinity Edit	✓	✓	✓		
	Organization: Edit Association Settings	✓				
	Organization: Edit Federation 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	✓				

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	Role: Create, Edit, Delete, or Copy	✓				
	Service Library: View service libraries	✓				
✓	SSL: Test Connection	✓	✓			
	UI Plugins: View	✓	✓	✓	✓	
✓ (Available in version 10.2.1 and later)	Truststore: Manage	✓				
✓ (Available in version 10.2.1 and later)	Truststore: View	✓				
	UI Plugins: View	✓	✓	✓	✓	
	vApp Template / Media: Copy	✓	✓	✓		
	vApp Template / Media: Create / Upload	✓	✓			
	vApp Template / Media: Edit	✓	✓	✓		
	vApp Template / Media: View	✓	✓	✓	✓	
	vApp Template: Add to My Cloud	✓	✓	✓	✓	
	vApp Template: Change Owner	✓	✓			
	vApp Template: Download	✓	✓			
	vApp: Change Owner	✓				
	vApp: Copy	✓	✓	✓	✓	
	vApp: Create / Reconfigure	✓	✓	✓		
	vApp: Delete	✓	✓	✓	✓	
	vApp: Download	✓	✓	✓		
	vApp: Edit Properties	✓	✓	✓	✓	
	vApp: Edit VM Compute Policy	✓	✓	✓		
	vApp: Edit VM CPU	✓	✓	✓		
	vApp: Edit VM Hard Disk	✓	✓	✓		

New in this release	Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
	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: Configure Logging	✓				
✓	VDC Group: View	✓				
	VDC Template: Instantiate	✓				
	VDC Template: View	✓				

Managing Rights Bundles

As a system administrator, you can create rights bundles and publish them to one and more organizations in your cloud. You can edit and delete existing rights bundles. You can unpublish rights bundles from organizations in your cloud.

Create a Rights Bundle

You can group a set of rights as a rights bundle which you can publish to one or more organizations in your system.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Rights Bundles**.
- 3 Click **Add**.
- 4 Enter a name and, optionally, a description for the new rights bundle.
- 5 Select the rights that you want to associate with this bundle.

The rights are grouped in categories and subcategories for view or manage access to the object to which they relate.

You can select the rights individually, by view or manage by subcategory, or by view or manage globally.

Category	Description
Access Control	Contains rights for viewing and managing organizations, rights, roles, and users.
Administration	Contains rights for viewing and managing general and multisite setting.
Compute	Contains rights for viewing and managing organization and provider VDCs, vApps, organization VDC templates, and VM monitoring.
Extensions	Contains rights for viewing and managing VMware Cloud Director plug-ins and extensions.
Infrastructure	Contains rights for viewing and managing vSphere resources.
Libraries	Contains rights for viewing and managing catalogs and catalog items.
Networking	Contains rights for viewing and managing network resources.

- 6 Click **Save**.

What to do next

You can publish the newly created rights bundle to one or more organizations in your system. See [Publish or Unpublish a Rights Bundle](#).

Clone a Rights Bundle

You can use an existing rights bundle as a template for the creation of a new bundle.

Prerequisites

Verify that you have the rights to add new roles to VMware Cloud Director.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Rights Bundles**.
- 3 Select the rights bundle that you want to clone and click **Clone**.
- 4 In the **Clone Rights Bundle** window, enter a name and description for the cloned bundle.
- 5 (Optional) To edit the cloned rights, turn on the **Modify Selected Rights** toggle, and select or deselect the rights you want to change for the cloned role.
- 6 Click **Save**.

Publish or Unpublish a Rights Bundle

You can publish a rights bundle to one or more organizations in your system. After you publish a rights bundle to an organization, the rights in this bundle become part of the organization set of rights.

Organization rights can comprise multiple rights bundles, but the organization administrators and users see a flat set of rights that they can use to create and modify roles.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Rights Bundles**.
- 3 Select the radio button next to the target bundle and click **Publish**.
- 4 To publish the bundle:
 - a Select **Publish to Tenants**.
 - b Select the organizations to which you want to publish the role.
 - If you want to publish the bundle to all existing and newly created organizations in your system, select **Publish to All Tenants**.
 - If you want to publish the bundle to particular organizations in your system, select the organizations individually.
- 5 To unpublish the bundle:
 - If you want to unpublish the bundle from all organizations in your system, deselect **Publish to Tenants**.
 - If you want to unpublish the bundle from particular organizations in your system, deselect **Publish to All Tenants**, and deselect the organizations individually.
- 6 Click **Save**.

Results

The rights in the published bundle are available in the selected organizations and can be used in the roles in these organizations.

The rights in the unpublished role are removed from the selected organizations and cannot be used in the roles in these organizations.

View and Edit a Rights Bundle

You can view the rights that are included in a rights bundle. You can modify the name, the description, and the rights of a bundle.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Rights Bundles**.
- 3 Click the name of the target bundle.

You can view the rights that are associated with the bundle by expanding the right categories.

- 4 Edit the bundle and click **Keep**.

Results

If you modified the rights of the bundle, the new set of rights is applied to all organizations to which this rights bundle is published.

Delete a Rights Bundle

You can remove a rights bundle that you no longer use in your organizations.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Rights Bundles**.
- 3 Select the radio button next to the target bundle and click **Delete**.
- 4 To confirm, click **OK**.

Managing Global Tenant Roles

As a system administrator, you can create global tenant roles and publish them to one or more organizations in your cloud. You can edit and delete existing global tenant roles. You can unpublish global tenant roles from individual organizations in your cloud.

After the initial VMware Cloud Director installation and setup, the system contains a set of predefined global tenant that are published to all organizations. See [Predefined Roles and Their Rights](#).

Create a Global Tenant Role

You can create a global tenant role that you can publish to one or more organizations in your system.

After the initial VMware Cloud Director installation and setup, the system contains predefined global tenant roles that are published to all organizations. For information about the predefined roles, see [Predefined Roles and Their Rights](#).

You can add custom global roles to your system.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Global Roles**.
- 3 Click **Add**.
- 4 Enter a name and, optionally, a description for the new role.
- 5 Select the rights that you want to associate with the role.

The rights are grouped in categories and subcategories for view or manage access to the object to which they relate.

You can select the rights individually, by view or manage by subcategory, or by view or manage globally.

Category	Description
Access Control	Contains rights for viewing and managing organizations, rights, roles, and users.
Administration	Contains rights for viewing and managing general and multisite setting.
Compute	Contains rights for viewing and managing organization and provider VDCs, vApps, organization VDC templates, and VM monitoring.
Extensions	Contains rights for viewing and managing VMware Cloud Director plug-ins and extensions.
Infrastructure	Contains rights for viewing and managing vSphere resources.
Libraries	Contains rights for viewing and managing catalogs and catalog items.
Networking	Contains rights for viewing and managing network resources.

- 6 Click **Keep**.

Results

Upon its creation, the new global tenant right is available only to the VMware Cloud Director Provider organization.

What to do next

You can publish the newly created role to one or more organizations in your system. See [Publish or Unpublish a Global Tenant Role](#).

Clone a Global Tenant Role

You can use an existing global tenant role as a template for the creation of a new role.

Prerequisites

Verify that you have the rights to add new roles to VMware Cloud Director.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Global Roles**.
- 3 Select the role that you want to clone and click **Clone**.
- 4 In the **Clone Global Role** window, enter a name and description for the cloned role.
- 5 (Optional) To edit the cloned rights, turn on the **Modify Selected Rights** toggle, and select or deselect the rights you want to change for the cloned role.
- 6 Click **Save**.

Publish or Unpublish a Global Tenant Role

You can publish a global tenant role to one or more organizations in your system. After you publish a role to an organization, this role becomes a part of the organization set of tenant roles.

Prerequisites

If you want to unpublish a global tenant role from an organization, verify that no user is assigned with this role in the organization.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Global Roles**.
- 3 Select the radio button next to the target role and click **Publish**.
- 4 To publish the role:
 - a Select **Publish to Tenants**.
 - b Select the organizations to which you want to publish the role.
 - If you want to publish the role to all existing and newly created organizations in your system, select **Publish to All Tenants**.
 - If you want to publish the role to particular organizations in your system, select the organizations individually.

- 5 To unpublish the role:
 - If you want to unpublish the role from all organizations in your system, deselect **Publish to Tenants**.
 - If you want to unpublish the role from particular organizations in your system, deselect **Publish to All Tenants**, and deselect the organizations individually.
- 6 Click **Save**.

Results

The published role is available in the selected organizations and can be assigned to users in these organizations. Organization administrators cannot edit global tenant roles that are published to their organizations.

The unpublished role is removed from the selected organizations and cannot be assigned to users in these organizations.

View and Edit a Global Tenant Role

You can view the rights that are included in a global tenant role. You can modify the name, the description, and the rights of a global tenant role.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, select **Global Roles**.
- 3 Click the name of the target role.

You can view the rights that are associated with the role by expanding the right categories.
- 4 To modify the name, the description, or the rights of the role, click **Edit**.
- 5 Edit the role and click **Keep**.

Results

If you modified the rights of the role, the new set of rights is applied to the users across all organizations that are assigned with this role.

Delete a Global Tenant Role

You can remove a global tenant role that you no longer use in your organizations.

Prerequisites

The global tenant role that you want to delete must not be assigned to any user across all organizations.

Procedure

- 1 From the top navigation bar, select **Administration**.

- 2 In the left panel, under **Tenant Access Control**, select **Global Roles**.
- 3 Select the radio button next to the target role and click **Delete**.
- 4 To confirm, click **OK**.

Managing Provider Roles

You can create and manage roles in your VMware Cloud Director Provider organization.

For information about managing tenant roles, see the *VMware Cloud Director Tenant Portal Guide*.

Create a Provider Role

You can create a role in your VMware Cloud Director Provider organization.

After the initial VMware Cloud Director installation and setup, the system contains predefined roles that are local to the Provider organization and global to all organizations. For information about the predefined roles, see [Predefined Roles and Their Rights](#).

You can add custom provider roles to your Provider organization.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Roles**.
- 3 Click **New**.
- 4 Enter a name and, optionally, a description for the new role.
- 5 Select the rights that you want to associate with the role.

The rights are grouped in categories and subcategories for view or manage access to the object to which they relate.

You can select the rights individually, by view or manage by subcategory, or by view or manage globally.

Category	Description
Access Control	Contains rights for viewing and managing organizations, rights, roles, and users.
Administration	Contains rights for viewing and managing general and multisite setting.
Compute	Contains rights for viewing and managing organization and provider VDCs, vApps, organization VDC templates, and VM monitoring.
Extensions	Contains rights for viewing and managing VMware Cloud Director plug-ins and extensions.
Infrastructure	Contains rights for viewing and managing vSphere resources.

Category	Description
Libraries	Contains rights for viewing and managing catalogs and catalog items.
Networking	Contains rights for viewing and managing network resources.

6 Click **Save**.

Results

The newly created role is available for assigning to users in your Provider organization.

Clone a Provider Role

You can use an existing provider role as a template for the creation of a new role.

Prerequisites

Verify that you have the rights to add new roles to VMware Cloud Director.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Roles**.
- 3 Select the role that you want to clone and click **Clone**.
- 4 In the **Clone Role** window, enter a name and description for the cloned role.
- 5 (Optional) To edit the cloned rights, turn on the **Modify Selected Rights** toggle, and select or deselect the rights you want to change for the cloned role.
- 6 Click **Save**.

View or Edit a Provider Role

You can view the rights that are included in a role that is local to your VMware Cloud Director Provider organization. You can modify the name, the description, and the rights of a role.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Roles**.
- 3 Click the name of the target role.
You can view the rights that are associated with the role by expanding the right categories.
- 4 To modify the name, the description, or the rights of the role, click **Edit**.
- 5 Edit the role and click **Save**.

Results

If you modified the rights of the role, the new set of rights is applied to the users that are assigned with this role.

Delete a Provider Role

You can remove a role that you no longer use in your VMware Cloud Director Provider organization.

Prerequisites

The role that you want to delete must not be assigned to any user.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Roles**.
- 3 Select the radio button next to the target role and click **Delete**.
- 4 To confirm, click **OK**.

Managing Provider Users and Groups

You can add and import users and groups to your VMware Cloud Director Provider organization.

For information about managing organization users and groups, see the *VMware Cloud Director Tenant Portal Guide*.

Managing Provider Users

You can manage the users in your Provider organization by using the Service Provider Admin Portal.

For information about managing tenant users in organizations, see the *VMware Cloud Director Tenant Portal Guide*.

Create a Provider User

You can create a user in your VMware Cloud Director Provider organization.

During the VMware Cloud Director installation and setup, you create a **system administrator** account. After the initial setup, you can create additional administrators and users to the Provider organization.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Users**.
- 3 Click **New**.

- 4 Enter a user name and password for the new user.

The password must contain at least six characters.

- 5 Select whether to enable the user upon creation.

- 6 From the **Available roles** drop-down menu, select a role for the user.

The list of available roles comprises the global roles and the roles that are local to your system organization.

- 7 (Optional) Enter contact information for the user.

You can enter the full name, email address, phone number, and instant messaging ID.

- 8 (Optional) Set the quotas for the user.

- a You can set a limit of the virtual machines owned by the user, or select **Unlimited**.

- b You can set a limit of the running virtual machines owned by the user, or select **Unlimited**.

Import Provider Users

You can import users to your VMware Cloud Director Provider organization from a previously configured LDAP or SAML identity provider.

Prerequisites

[Configure a System LDAP Connection](#) or [Configure Your System to Use a SAML Identity Provider](#).

Procedure

- 1 From the top navigation bar, select **Administration**.

- 2 In the left panel, under **Provider Access Control**, select **Users**.

- 3 Click **Import Users**.

- 4 From the **Source** drop-down menu, select your identity provider type.

Can be **LDAP** or **SAML**.

If you configured only one identity provider, this option is hard-coded.

5 Specify the users.

Option	Description
LDAP	<ol style="list-style-type: none"> Enter a full or partial name of a user and click Search. From the search results, select the users that you want to import. From the Assign Role drop-down menu, select a role for the imported users.
SAML	<ol style="list-style-type: none"> Enter the user names of the users that you want to import in the name identifier format supported by the SAML identity provider. Use a new line for each user name. From the Assign Role drop-down menu, select a role for the imported users.

6 Click **Save**.

Results

You can see the imported users in the list of users.

Edit a Provider User

You can change the password, role, contact information, and quotas of a user in your Provider organization. You cannot change the user name.

Procedure

- From the top navigation bar, select **Administration**.
- In the left panel, under **Provider Access Control**, select **Users**.
- Click the radio button next to the name of the target user and click **Edit**.
- Edit the user details and click **Save**.

Activate or Deactivate a Provider User

After you deactivate a user, the user cannot log in to VMware Cloud Director.

Procedure

- From the top navigation bar, select **Administration**.
- In the left panel, under **Provider Access Control**, select **Users**.
- Click the radio button next to the name of the target user and click **Disable** or **Enable**.
- If deactivating a user, click **OK** to confirm.

Delete a Provider User

You can remove a user from your VMware Cloud Director Provider organization by deleting the user account.

To delete a stranded user that lost access to the system because their LDAP group was deleted, use the VMware Cloud Director API.

Prerequisites

Deactivate the user that you want to delete. See [Activate or Deactivate a Provider User](#).

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Users**.
- 3 Click the radio button next to the name of the target user and click **Delete**.
- 4 To confirm, click **OK**.

Unlock a Provider User

If you enabled account lockout in your password policy system settings, users might lock their accounts after a certain number of invalid login attempts. Even if the lockout is set with an account lockout interval, you can unlock a user account without waiting for the lock to expire.

For information about configuring the account lockout policy, see [Configure the Password Policy](#).

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Users**.
- 3 Click the radio button next to the name of the target user and click **Unlock**.

Managing Provider Groups

You can import, edit, and delete groups from your Provider organization by using the Service Provider Admin Portal.

For information about managing groups in organizations, see the *VMware Cloud Director Tenant Portal Guide*.

Import a Provider Group

You can import groups to your VMware Cloud Director Provider organization from a previously configured LDAP or SAML identity provider.

Prerequisites

[Configure a System LDAP Connection](#) or [Configure Your System to Use a SAML Identity Provider](#).

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Groups**.

- 3 Click **Import Groups**.
- 4 From the **Source** drop-down menu, select your identity provider type.
Can be **LDAP** or **SAML**.
If you configured only one identity provider, this option is hard-coded.
- 5 Specify the users.

Option	Description
LDAP	<ol style="list-style-type: none"> a Enter a full or partial name of a group and click Search. b From the search results, select the groups that you want to import. c From the Assign Role drop-down menu, select a role for the users in the imported groups.
SAML	<ol style="list-style-type: none"> a Enter the names of the groups that you want to import in the name identifier format supported by the SAML identity provider. Use a new line for each group name. b From the Assign Role drop-down menu, select a role for the users in the imported groups.

- 6 Click **Save**.

Edit a Provider Group

You can edit the description and change the role of the members of a group that you previously imported to your VMware Cloud Director Provider organization.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Groups**.
- 3 Click the radio button next to the name of the target group and click **Edit**.
- 4 Edit the group details, and click **Save**.

Delete a Provider Group

You can remove a group from your VMware Cloud Director Provider organization

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Provider Access Control**, select **Groups**.
- 3 Click the radio button next to the name of the target group and click **Delete**.
- 4 To confirm, click **OK**.

Managing System Settings

11

A VMware Cloud Director system administrator can control system-wide settings related to LDAP, email notification, licensing, and general system preferences.

This chapter includes the following topics:

- [Modify General System Settings](#)
- [General System Settings](#)
- [Activate FIPS Mode on the Cells in the Server Group](#)
- [Configure the System Email Settings](#)
- [Change the VMware Cloud Director License](#)
- [Configure the Catalog Synchronization Settings](#)
- [Create an Advisory Dashboard](#)
- [Configuring and Monitoring Blocking Tasks and Notifications](#)
- [Configure Public Addresses](#)
- [Managing Identity Providers](#)
- [Managing Certificates](#)
- [Managing Plug-Ins](#)
- [Customizing the VMware Cloud Director Portals](#)
- [Configure the Password Policy](#)
- [Configure vSphere Services](#)

Modify General System Settings

VMware Cloud Director includes general system settings related to activity logs, networking, session timeouts, certificates, organization limits, operation limits, and so on. The default settings are appropriate for many environments, but you can modify the settings to meet your needs.

For a list of the properties that you can modify, see [General System Settings](#).

Note For information about changing the date, time, or time zone of the VMware Cloud Director appliance, see <https://kb.vmware.com/kb/59674>.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Settings**, click **General**.
- 3 Click **Edit** for the section you want to modify, edit the properties, and click **Save**.

General System Settings

VMware Cloud Director includes general system settings that you can modify to meet your needs.

Table 11-1. General System Settings

Name	Category	Description
Activity log history to keep	Activity Log	Number of days of the log history to keep before deleting it. Enter 0 never to delete logs.
Activity log history shown	Activity Log	Number of days of the log history to display. To show all activity, enter 0 .
Display debug information	Activity Log	Enable this setting to display the debug information in the VMware Cloud Director task log.
IP address release timeout	Networking	Number of seconds to keep released IP addresses on hold before making them available for allocation again. This default setting is 2 hours (7200 seconds) to allow old entries to expire from client ARP tables.
Allow Overlapping External Networks	Networking	To add external networks that run on the same network segment, select the check box. Enable this setting only if you are using non-VLAN-based methods to isolate your external networks.
Allow FIPS mode	Networking	Allows enablement of FIPS mode on Edge Gateways. Requires NSX 6.3 or later. See FIPS Mode in the <i>VMware NSX for vSphere</i> documentation.
Default syslog server settings for networks	Networking	Enter IP addresses for up to two Syslog servers for networks to use. This setting does not apply to Syslog servers used by cloud cells.
Provider Locale	Localization	Select a locale for provider activity, including log entries, email alerts, and so on.
Idle session timeout	Timeouts	Amount of time the VMware Cloud Director application remains active without a user interaction.
Maximum session timeout	Timeouts	Maximum amount of time the VMware Cloud Director application remains active.

Table 11-1. General System Settings (continued)

Name	Category	Description
Host refresh frequency	Timeouts	How often VMware Cloud Director checks whether its ESXi hosts are accessible or inaccessible.
Host hung timeout	Timeouts	Select the amount of time to wait before marking a host as hung.
Transfer session timeout	Timeouts	Amount of time to wait before failing a paused or canceled upload task, for example upload media or upload vApp template. This timeout does not affect upload tasks that are in progress.
Enable upload quarantine with a timeout of __ seconds	Timeouts	Select the check box and enter a timeout number representing the amount of time to quarantine uploaded files.
Verify vCenter and vSphere SSO certificates	Certificates	VMware Cloud Director always verifies the certificates. When enabled, verifies the host names in the vCenter Server certificates.
Verify NSX Manager certificates	Certificates	VMware Cloud Director always verifies the certificates. When enabled, VMware Cloud Director verifies the host names in the NSX Manager certificates.
Edit Organization Limits	Organization VDC Limits	Enter the maximum number of organization virtual data centers per organization, or select Unlimited .
Number of resource intensive operations running per user	Operation Limits	Enter the maximum number of simultaneous resource-intensive operations per user, or select Unlimited .
Number of resource intensive operations to be queued per user (in addition to running)	Operation Limits	Enter the maximum number of queued resource-intensive operations per user, or select Unlimited .
Number of resource intensive operations running per organization	Operation Limits	Enter the maximum number of simultaneous resource-intensive operations per organization, or select Unlimited .
Number of resource intensive operations to be queued per organization	Operation Limits	Enter the maximum number of queued resource-intensive operations per organization, or select Unlimited .
Provide default vApp names	Other	Select the check box to configure VMware Cloud Director to provide default names for new vApps.
Make Allocation pool Org VDCs elastic	Other	Select the check box to enable the elastic allocation pool, making all allocation pool organization virtual data centers elastic. Before deselecting this option, ensure all virtual machines for each organization virtual data center have been migrated to a single cluster.
VM discovery enabled	Other	By default, each organization VDC automatically discovers vCenter VMs that were created in any resource pool that backs the VDC. Clear to deactivate this setting for all VDCs in the system.

Activate FIPS Mode on the Cells in the Server Group

You can configure VMware Cloud Director 10.2.2 and later on Linux to use FIPS 140-2 validated cryptographic modules and to run in FIPS-compliant mode.

The Federal Information Processing Standard (FIPS) 140-2 is a U.S. and Canadian government standard that specifies security requirements for cryptographic modules. The NIST Cryptographic Module Validation Program (CMVP) validates the cryptographic modules compliant with the FIPS 140-2 standards.

The goal of VMware Cloud Director FIPS support is to ease the compliance and security activities in various regulated environments. To learn more about support for FIPS 140-2 in VMware products, see <https://www.vmware.com/security/certifications/fips.html>.

In VMware Cloud Director, FIPS-validated cryptography is deactivated by default. By activating FIPS mode, you configure VMware Cloud Director to use FIPS 140-2 validated cryptographic modules and to run in FIPS-compliant mode.

Note Activating FIPS mode also activates reverse lookup of host names.

Important When you activate FIPS mode, the integration with vRealize Orchestrator does not work.

In VMware Cloud Director 10.2.2 when you activate FIPS mode, you cannot encrypt SAML assertions. When not in FIPS mode, there is no restriction on assertion encryption.

VMware Cloud Director uses the following FIPS 140-2 validated cryptographic modules:

- VMware's BC-FJA (Bouncy Castle FIPS Java API), version 1.0.2.1: [Certificate #3673](#)
- VMware's OpenSSL FIPS Object Module, version 2.0.20-vmw: [Certificate #3857](#)

VMware Cloud Director is in a bundle with the cell management tool (CMT). However, the cell management tool is not FIPS-compliant.

For information about activating FIPS mode on the VMware Cloud Director appliance, see [Activate or Deactivate FIPS Mode on the VMware Cloud Director Appliance](#).

Prerequisites

- Verify that the certificates have the `KeyCertSign` bit asserted by using OpenSSL. FIPS mode can function only if the VMware Cloud Director SSL certificates have the `KeyCertSign` asserted.

```
openssl crl2pkcs7 -nocrl -certfile certificates.pem | openssl pkcs7 -print_certs -text -noout
```

If the certificates do not include the extension, specify the `KeyCertSign` bit when creating an SSL certificate keystore.

- Install and activate the `rng-tools` set of utilities. See <https://wiki.archlinux.org/index.php/Rng-tools>.

- If metrics collection is activated, verify that the Cassandra certificates follow the X.509 v3 certificate standard and include all the necessary extensions. You must configure Cassandra with the same cipher suites that VMware Cloud Director uses. For information about the allowed SSL ciphers, see [Managing the List of Allowed SSL Ciphers](#).
- Unregister VMware Cloud Director from the vCenter Lookup Service. See [Configure vSphere Services](#).

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Settings**, select **SSL**.
- 3 Click **Enable**.
- 4 Confirm that your environment meets all prerequisites to activating FIPS mode.
If your environment does not meet all prerequisites before starting the FIPS mode configuration, VMware Cloud Director might become inaccessible.
- 5 To confirm you want to start the process, click **Enable**.
When the configuration finishes, VMware Cloud Director displays a message to restart your cloud cells.
- 6 After VMware Cloud Director displays a message to restart your cloud cells, restart every cell in the VMware Cloud Director server group.

What to do next

- Deactivate FIPS mode by clicking **Disable**, and after VMware Cloud Director indicates that the configuration is ready, restart the cells.
- You can view the FIPS status of the active VMware Cloud Director cells by using the `fips-mode` CMT command. See [View the FIPS Status of All Active Cells](#) in the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

Configure the System Email Settings

You can edit the system email settings, including configuring the SMTP server settings and VMware Cloud Director notification settings.

VMware Cloud Director requires an SMTP server to send user notifications and system alert emails to system users.

VMware Cloud Director sends system alert emails when it has important information to report. For example, VMware Cloud Director sends an alert when a datastore is running out of space. You can configure VMware Cloud Director to send email alerts to all system administrators or to a specified list of email addresses.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 From the left pane, under **Settings**, select **Email**, and click **Edit**.
- 3 Enter the DNS host name or IP address of the SMTP mail server.
- 4 Enter the SMTP server port number.
- 5 (Optional) If the SMTP server requires a user name, toggle on the **Requires authentication** option and enter the user name and password for the SMTP account.
- 6 Select the **Notification Settings** tab.
- 7 Enter an email address to appear as the sender for VMware Cloud Director emails.

VMware Cloud Director uses the sender's email address to send runtime and storage lease expiration alerts.
- 8 (Optional) Enter text for the subject prefix.
- 9 Select the recipients of the notifications.

By default, only organization administrators receive the SMTP notifications.
- 10 Click **Save**.
- 11 (Optional) Test the SMTP settings.
 - a Click **Test**.
 - b If you enabled the **Requires authentication** option, enter the SMTP server password.
 - c Enter a destination email address and click **Test**.

Change the VMware Cloud Director License

VMware Cloud Director requires a valid license, specified as a serial number, to run. You can modify the licensing information that you entered during the initial VMware Cloud Director configuration.

The VMware Cloud Director product serial number is not the same as your vCenter Server license key. You can obtain a VMware Cloud Director serial number from the VMware License Portal.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 From the left pane, select **License** and click **Edit**.
- 3 Enter a new serial number and click **Save**.

Configure the Catalog Synchronization Settings

You can edit the catalog synchronization settings for all organizations and catalogs, including the refresh rate of the catalog subscriptions.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 From the left pane, under **Settings**, select **Catalog**.
- 3 Click **Edit**.
- 4 Enable the catalog synchronization.
- 5 Set the synchronization start and stop times.
- 6 Set the synchronization interval.

The synchronization interval is the refresh rate of the catalog subscriptions.

- 7 Click **Save**.

What to do next

For information about configuring catalog synchronization throttling, see the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

Create an Advisory Dashboard

You can create notifications that appear on top of the UI pages in the VMware Cloud Director Service Provider Admin Portal and the Tenant Portal. The messages can appear to system administrators, the users within an organization, or the users in all organizations.

You cannot edit advisories once you create them.

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.

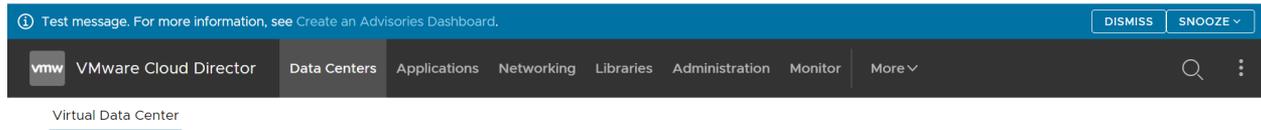
- 5 Select the period for which you want the notification to appear in the UI.

You can view all advisories in the **Advisories** tab, however they appear to the selected group of users only during the selected period.

- 6 Select whether you want the notification to appear only to system administrators, to all users within the organization or across organizations.
- 7 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.

Configuring and Monitoring Blocking Tasks and Notifications

You can use blocking tasks and notifications to configure VMware Cloud Director to send AMQP messages triggered by certain events.

Some of these messages are simply notifications that the event has occurred. Other messages publish information to a designated AMQP endpoint indicating that a requested action has been blocked and is pending action by a client application bound to that endpoint. These messages are known as blocking tasks.

A **system administrator** can configure a system-wide set of blocking tasks that are subject to a programmatic action by an AMQP client.

Configure an AMQP Broker

If you want VMware Cloud Director to send AMQP messages triggered by certain events, you must configure an AMQP broker. You can use the AMQP messages to automate the handling of an underlying user request.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 Under **Settings**, select **Extensibility**.
The **AMQP Broker** tab opens.
- 3 Click the **Edit** button of the **AMQP Broker** section.
- 4 Enter the DNS host name or IP address of the AMQP host.
The fully qualified domain name of the RabbitMQ server host, for example, *amqp.example.com*.

- 5 Enter the AMQP port.

The default port at which the broker listens to messages is 5672.

- 6 Enter the exchange.

- 7 Enter the vHost.

The default is /.

- 8 Enter the prefix.

- 9 (Optional) To use SSL, turn on the **Use SSL** toggle and select one of the certificate options.

By default, the VMware Cloud Director AMQP service sends unencrypted messages. You can configure the AMQP service to encrypt these messages by using SSL. You can also configure the service to verify the broker certificate by using the default JCEKS trust store of the Java runtime environment on the VMware Cloud Director cell, typically at `$VCLLOUD_HOME/jre/lib/security/cacerts`.

Option	Description
Accept all certificates	The CN record from the certificate owner field must match the AMQP broker host name. To use certificates that do not match the broker host name, turn on the Accept all certificates toggle.
SSL Certificate	Upload the SSL certificate.
SSL Key Store (JCEKS)	Upload the SSL keystore and enter the keystore password.

- 10 Enter a user name and password to connect to the AMQP host.

- 11 Click **Save**.

- 12 (Optional) To test the settings, click the **Test** button under the **AMQP Broker** section and provide the password.

- 13 (Optional) To publish audit events to the AMQP broker, click the **Edit** button under the **Non-blocking AMQP Notifications** section and turn on the **Enable notifications** toggle.

Configure Blocking Task Settings

You can configure certain operations as blocking tasks. These operations are suspended until a **system administrator** acts on them or a preconfigured timer expires. You can specify the timeout settings and default actions for blocking tasks. The settings apply to all organizations in the installation.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 Under **Settings**, select **Extensibility**.
- 3 Select the **Blocking Tasks** tab.

- 4 To edit the default extension timeout and default timeout action, click the **Edit** button under the **General** section.
 - a Edit the **Default blocking task timeout**.
 - b Edit the **Default Timeout Action**.

The **Default Timeout Action** is the action after a **Default blocking task timeout** expires.
 - c Click **Save**.
- 5 To edit the list of operations, considered as blocking tasks, click the **Edit** button under the **Operations** section.
 - a Select or deselect operations from the list of blocking tasks.
 - b Click **Save**.

Monitor Blocked Tasks

You can monitor the current blocked tasks or manually cancel, fail, or resume the tasks before the preconfigured timer expires.

Prerequisites

[Configure Blocking Task Settings](#)

Procedure

- 1 From the top navigation bar, under **Monitor**, select **Blocking Tasks**.

The tab displays a list of the current blocked tasks.
- 2 Select the task that you want to edit manually.
- 3 Decide between canceling, failing, or resuming the task and click the corresponding button.
- 4 Enter a message and click **Save**.

The message appears in the task details.

Configure Public Addresses

To fulfill load balancer or proxy requirements, you can change the default endpoint Web addresses for the VMware Cloud Director Web Portal, VMware Cloud Director API, and console proxy.

Public addresses are Web addresses exposed to clients of VMware Cloud Director. Defaults for these addresses are specified during installation. If necessary, you can update the addresses.

If VMware Cloud Director consists of a single cell, the installer creates public endpoints that usually provide sufficient access for API and Web clients. Installations and deployments that include multiple cells typically place a load balancer between the cells and the clients. Clients access the system at the load balancer's address. The load balancer distributes client requests across the available cells. Other network configurations that include a proxy or place the cells in a DMZ also require customized endpoints. Endpoint URL details are specific to your network configuration.

The endpoints for the VMware Cloud Director Tenant Portal and VMware Cloud Director Web Console require SSL certificates, preferably signed. You must specify a path to these certificates when you install or deploy VMware Cloud Director. If you customize any of these endpoints after installation or deployment, you might need to install new certificates that match endpoint details such as `hostname` and `subject alternative name`.

For the VMware Cloud Director appliance, you must configure the VMware Cloud Director public console proxy address, because the appliance uses a single IP address with custom port 8443 for the console proxy service. See [Step 6](#).

Prerequisites

Verify that you are logged in as a **system administrator**. Only a **system administrator** can customize the public endpoints.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Settings**, click **Public Addresses**.
- 3 To customize the public endpoints, click **Edit**.
- 4 To customize the VMware Cloud Director URLs, edit the **Web Portal** endpoints.
 - a Enter a custom VMware Cloud Director public URL for HTTP (non-secure) connections.
 - b Enter a custom VMware Cloud Director public URL for HTTPS (secure) connections and click **Upload** to upload the certificates that establish the trust chain for that endpoint.

The certificate chain must match the certificate used by the service endpoint, which is the certificate uploaded to each VMware Cloud Director cell keystore with alias `consoleproxy`. SSL termination of console proxy connections at a load balancer is not supported. The certificate chain must include an endpoint certificate, intermediate certificates, and a root certificate in the `PEM` format without a private key.

5 (Optional) To customize the Cloud Director REST API and OpenAPI URLs, turn off the **Use Web Portal Settings** toggle.

a Enter a custom HTTP base URL.

For example, if you set the HTTP base URL to **http://vcloud.example.com**, you can access the VMware Cloud Director API at `http://vcloud.example.com/api`, and you can access the VMware Cloud Director OpenAPI at `http://vcloud.example.com/cloudapi`.

b Enter a custom HTTPS REST API base URL and click **Upload** to upload the certificates that establish the trust chain for that endpoint.

For example, if you set the HTTPS REST API base URL to **https://vcloud.example.com**, you can access the VMware Cloud Director API at `https://vcloud.example.com/api`, and you can access the VMware Cloud Director OpenAPI at `https://vcloud.example.com/cloudapi`.

The certificate chain must match the certificate used by the service endpoint, which is either the certificate uploaded to each VMware Cloud Director cell keystore with alias `http` or the load balancer VIP certificate if an SSL termination is used. The certificate chain must include an endpoint certificate, intermediate certificates, and a root certificate in the PEM format without a private key.

6 Enter a custom VMware Cloud Director public console proxy address.

■ Customize the VMware Cloud Director appliance public console proxy address.

This address is the fully qualified domain name (FQDN) of the VMware Cloud Director appliance `eth0` NIC, specified either by FQDN or IP address, with custom port `8443` for the console proxy service.

■ Customize the VMware Cloud Director on Linux public console proxy address.

This address is the fully qualified domain name (FQDN) of the VMware Cloud Director server or load-balancer with the port number. The default port is `443`.

For example, for a VMware Cloud Director appliance instance with FQDN `vcloud.example.com`, enter **vcloud.example.com:8443**.

VMware Cloud Director uses the console proxy address when opening a remote console window on a VM.

7 Click **Save**.

Managing Identity Providers

You can integrate your cloud with an external identity provider and import users and groups to your organizations. You can configure an LDAP server connection at a system or organization level. You can configure a SAML integration at an organization level.

Managing LDAP Connections

As a system administrator, you can configure your VMware Cloud Director system organization and any other organization in the system to use an LDAP server as a source of users and groups. The organizations can use either the system LDAP connection or a private LDAP connection.

Starting with version 10.1, VMware Cloud Director is moving to a centralized, tenant-aware storage area for certificate management. This way, VMware Cloud Director centralizes all certificates in one place so that **system administrators** and **organization administrators** can view, audit, and manage all certificates in use by various components in the system. You can use the VMware Cloud Director API to add, update, or remove certificates from the new tenant-aware storage area. See *VMware Cloud Director API Schema Reference*.

When adding or editing a new LDAP server endpoint, the VMware Cloud Director UI probes that endpoint for any certificates it is presenting. VMware Cloud Director adds to a centralized certificate storage area any certificate you decide to trust.

Configure a System LDAP Connection

To provide VMware Cloud Director and its organizations with shared access to users and groups, you can configure an LDAP connection at a system level.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Identity Providers**, click **LDAP**.

The current LDAP settings are displayed.

What to do next

[Configure, Test, and Synchronize an LDAP Connection.](#)

Configure an Organization LDAP Connection

You can configure an organization to use the system LDAP connection as a shared source of users and groups. You can configure an organization to use a separate LDAP connection as a private source of users and groups.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Organizations**.
- 3 Click the name of the target organization.
You are redirected to the VMware Cloud Director Tenant Portal of the organization.
- 4 From the top navigation bar, select **Administration**.
- 5 In the left panel, under **Identity Providers**, click **LDAP**.

The current LDAP settings are displayed.

- 6 On the **LDAP Options** tab, click **Edit**.
- 7 Configure the LDAP source of users and groups for this 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.
VCD system LDAP service	The organization uses the VMware Cloud Director system LDAP connection that you previously configured. See Configure a System LDAP Connection .
Custom LDAP service	The organization uses a private LDAP server as a source of organization users and groups. Click the Custom LDAP tab and 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>.

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.

Required Information	Description
Base distinguished name	<p>The base distinguished name (DN) is the location in the LDAP directory where VMware Cloud Director to connect.</p> <p>To connect at root level, enter only the domain components, for example, DC=example,DC=com.</p> <p>To connect to a node in the domain tree structure, enter the distinguished name for that node, for example, OU=ServiceDirector,DC=example,DC=com.</p> <p>Connecting to a node limits the scope of the directory available to VMware Cloud Director.</p>
Connector type	The type of your LDAP server. Can be Active Directory or OpenLDAP .
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 Upload button and import an LDAP SSL certificate or select Accept all certificates .
Authentication method	<p>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.</p> <p>If you want to use Kerberos, you must configure the LDAP connection by using the vCloud API.</p>
User name	<p>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.</p> <p>If the anonymous read support is enabled on your LDAP server, you can leave these text boxes blank.</p>
Password	<p>The password for the service account that connects to the LDAP server.</p> <p>If the anonymous read support is enabled on your LDAP server, you can leave these text boxes blank.</p>

- 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.

Configure Your System to Use a SAML Identity Provider

If you want to import users and groups from a SAML identity provider to your system organization, you must configure your system organization with this SAML identity provider. Imported users can log in to the system organization with the credentials established in the SAML identity provider.

To configure VMware Cloud Director with a SAML identity provider, you establish a mutual trust by exchanging SAML service provider and identity provider metadata.

When an imported user attempts to log in, the system extracts the following attributes from the SAML token, if available, and use them for interpreting the corresponding pieces of information about the user.

- `email address = "EmailAddress"`
- `user name = "UserName"`
- `full name = "FullName"`
- `user's groups = "Groups"`
- `user's roles = "Roles"` (this attribute is configurable)

Group information is used if the user is not directly imported but is expected to log in by virtue of membership in imported groups. A user can belong to multiple groups, so 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 using API 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 has no any rights to perform any activities.

Tip If you need to log in as a local user, you can use the base URL that you configured, such as `https://vcloud.example.com/tenant/tenant_name/login`.

Prerequisites

- Verify that you have access to a SAML 2.0 compliant identity provider.
- Obtain an XML file with the following metadata from your SAML identity provider.
 - 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, consult the documentation for your SAML provider.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under Identity Providers, click **SAML** and click **Edit**.
The current SAML settings are displayed.
- 3 From the **Service Provider** tab, download the VMware Cloud Director SAML service provider metadata.
 - a Enter an Entity ID for the system organization.
The Entity ID uniquely identifies your system organization to your Identity Provider.
 - b Examine the certificate expiration date and, if expiring soon, regenerate the certificate by clicking **Regenerate**.
The certificate is included in the SAML metadata, and is used for both encryption and signing. Either or both of these might be required depending on how trust is established between your organization and your SAML IDP.
 - c Click the **Metadata** link.
The link is similar to `https://VCD_host_name/cloud/org/System/saml/metadata/alias/vcd`.
Your browser downloads the SAML service provider metadata, an XML file which you must provide to your identity provider.

- 4 On the **Identity Provider** tab, upload the SAML metadata that you previously received from your identity provider.
 - a Select **Use SAML Identity Provider**.
 - b Either click the **Browse** icon and upload the file, or copy and paste its content in the **Metadata XML** text box.
- 5 Click **Save**.

Managing Certificates

You can import, download, edit, and delete certificates from VMware Cloud Director. You can copy the certificate PEM data to the clipboard.

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](#).

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.

You must manage the certificate libraries separately for each site.

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](#).

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 Plug-Ins

VMware Cloud Director plug-ins extend the functions of the Service Provider Admin Portal and the VMware Cloud Director Tenant Portal. You can upload, deactivate, and delete plug-ins from the Service Provider Admin Portal. You can publish a plug-in to the service provider and individual organizations.

Some plug-ins are installed as part of VMware Cloud Director.

CPOM extension

Provides the capability for viewing and managing dedicated vCenter Server instances and proxies by using the VMware Cloud Director Tenant Portal.

Customize Portal

Provides the capability for customizing the VMware Cloud Director Service Provider Admin Portal and the VMware Cloud Director Tenant Portal.

vCloud Availability

The VMware vCloud[®] Availability[™] plug-in provides the capability to access vCloud Availability Portal directly from the VMware Cloud Director user interface. For more information, see [vCloud Availability Documentation](#).

Upload a Plug-in

You can upload additional plug-ins to your VMware Cloud Director Service Provider Admin Portal for use by the service provider and organizations in the cloud.

Prerequisites

Download the plug-in installation file.

Procedure

- 1 From the top navigation bar, select **More > Customize Portal**.
- 2 Click **Upload**.
- 3 Click **Select plugin file**, browse to the target installation file, and click **Open**.
- 4 Click **Next**.
- 5 Select the scope for this plug-in.

Option	Description
Service Providers	The plug-in function becomes available in the VMware Cloud Director Service Provider Admin Portal.
Tenants	The plug-in function becomes available in the VMware Cloud Director Service Provider Admin Portal of the organizations that you select.

- 6 If you scoped the plug-in to tenants, select the organizations to which you want to publish this plug-in.
- 7 Review the **Review & Finish** page, and click **Finish**.

Activate or Deactivate a Plug-in

To prevent all organizations from using a plug-in, you can deactivate this plug-in.

Procedure

- 1 From the top navigation bar, select **More > Customize Portal**.

- 2 Select the check box next to the names of the target plug-ins, and click **Enable** or **Disable**.

Delete a Plug-in

You can remove one or more plug-ins from the VMware Cloud Director Service Provider Admin Portal.

Procedure

- 1 From the top navigation bar, select **More > Customize Portal**.
- 2 Select the check boxes next to the names of the plug-ins that you want to remove, and click **Delete**.
- 3 To confirm, click **Save**.

Publish or Unpublish a Plug-in from an Organization

You can modify the set of organizations that can use the function provided by a plug-in.

You can modify the set of organizations for multiple plug-ins.

Procedure

- 1 From the top navigation bar, select **More > Customize Portal**.
- 2 Select the check boxes next to the names of the target plug-ins, and click **Publish**.
- 3 Select the scope for this plug-in.

Option	Description
Service Providers	The plug-in function becomes available in the VMware Cloud Director Service Provider Admin Portal.
Tenants	The plug-in function becomes available in the VMware Cloud Director Service Provider Admin Portal of the organizations that you select.

- 4 If you scoped the plug-in to tenants, select the organizations to which you want to publish this plug-in.
- 5 Click **Save**.

Customizing the VMware Cloud Director Portals

To match your corporate branding standards and to create a fully custom cloud experience, you can set the logo and the theme for your VMware Cloud Director Service Provider Admin Portal and for the VMware Cloud Director Tenant Portal of each organization. In addition, you can modify and add custom links to the two upper right menus in the VMware Cloud Director portals.

Note To customize your branding attributes and links, you must use the `branding` vCloud OpenAPI methods. See *Getting Started with VMware Cloud Director OpenAPI* at <https://code.vmware.com>.

Portal Branding

As part of the installation, VMware Cloud Director contains two themes - default and dark. You can create, manage, and apply custom themes. In addition, you can change the portal name, the logo, and the browser icon. In addition, the browser title adopts the portal name that you set.

You set the branding attributes at a system level, so that you customize the VMware Cloud Director Service Provider Admin Portal. The VMware Cloud Director Tenant Portal for each organization adopts the system branding attributes unless you configured branding attributes for the particular tenant.

For a particular tenant, you can selectively override any combination of the portal name, background color, logo, icon, theme, and custom links. Any value that you do not set uses the corresponding system default value.

Note By default, the individual tenant branding is not shown outside of a logged in session. The individual tenant branding does not appear on login and logout pages, so that tenants cannot discover the existence of other tenants. You can enable branding outside of logged in sessions by using the cell management tool:

```
manage-config -n backend.branding.requireAuthForBranding -v false
```

For information about using the cell management tool, see the *VMware Cloud Director Installation, Configuration, and Upgrade Guide*.

Custom Links

Custom links are a component of the portal branding. There are two types of custom links:

- `override` menu items replace the existing links for menu items **Help**, **About**, and **Download VMRC**. By default, **Download VMRC** redirects the users to <https://my.vmware.com> to download VMRC, which requires users to have registered accounts for downloading. By overriding this link, you can relocate the VMRC installer to your own server.
- `link` menu items are new links that you add to the **Log out** menu item in the upper right corner of the portal. The new custom links appear in the order given in the API call.

You can organize these custom links by using `section` and `separator` menu items. A `section` menu item adds a header to the menu, a `separator` menu item adds a line to the menu.

Custom links support custom variables which you can use to pass identifying information to other applications in the form of query parameters.

VMware Cloud Director supports the following custom variables in the `url` value for a custom link:

Table 11-2. Custom Variables for Custom Links

Variable	Description
<code>\${TENANT_NAME}</code>	Organization name
<code>\${TENANT_ID}</code>	Organization ID
<code>\${SESSION_TOKEN}</code>	x-vcloud-authorization token

For example,

```
url: https://host:port/tenant/${TENANT_NAME}/vdc
```

in the VMware Cloud Director Tenant Portal for organization myorg is converted to:

```
url: https://host:port/tenant/myorg/vdc
```

Configure the Password Policy

To prevent a user from logging in to VMware Cloud Director after a certain number of failed attempts, you can enable the account lockout.

Changes to the system account lockout policy apply to all new organizations. Organizations created before the account lockout policy change must be changed at the organization level.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 In the left panel, under **Settings**, click **Password Policy**.
- 3 Click **Edit**.
- 4 To enable the account lockout, turn on the **Account lockout** toggle.
- 5 Select the accepted number of invalid logins before locking an account.
- 6 Select the lockout interval.
- 7 To enable the **system administrator** account lockout, turn on the **System Administrator account can be locked out** toggle.
- 8 Click **Save**.

Configure vSphere Services

You can configure and enable VMware Cloud Director to use vCenter Single Sign-On so that the vSphere identity provider authenticates the system administrators.

vCenter Lookup Service contains topology information about the vSphere infrastructure, enabling vSphere components to connect to each other securely.

Procedure

- 1 From the top navigation bar, select **Administration**.
- 2 From the left pane, under **Settings**, select **vSphere Services**.
- 3 Configure the vSphere Services.
 - To register VMware Cloud Director with the vCenter Lookup Service, click **Register**.
 - To unregister VMware Cloud Director from the vCenter Lookup Service, click **Unregister**.
- 4 Enter the vCenter Lookup Service URL, for example, `https://hostname:443/lookupservice/sdk`.
- 5 Enter the user name and password of a vCenter Single Sign-On user with administrative privileges, for example, the `administrator@your_domain_name` user.

Results

If you registered VMware Cloud Director with the vCenter Lookup Service, **system administrators** must log in to VMware Cloud Director with their vCenter Single Sign-On credentials.

Monitoring VMware Cloud Director

12

System administrators can monitor completed and in-progress operations and view resource usage information at the provider virtual data center, organization virtual data center, and datastore level.

Starting with version 9.1, VMware Cloud Director does not support VMware vCenter Chargeback Manager. See the [VMware Product Interoperability Matrices](#).

This chapter includes the following topics:

- [VMware Cloud Director and Cost Reporting](#)
- [View Use Information for a Provider Virtual Data Center](#)

VMware Cloud Director and Cost Reporting

You can use VMware vRealize Operations Tenant App for VMware Cloud Director to configure a cost reporting system for VMware Cloud Director.

The VMware vRealize Operations Tenant App features metering capabilities that allow service providers to provide their customer base with chargeback services.

The VMware vRealize Operations Tenant App is also a tenant facing application which provides tenant administrators with visibility to their environment and to their billing data.

For information about compatibility between VMware Cloud Director and VMware vRealize Operations Tenant App, see the *VMware Product Interoperability Matrixes* at http://partnerweb.vmware.com/comp_guide/sim/interop_matrix.php.

You can download the VMware vRealize Operations Tenant App at <https://marketplace.vmware.com/vsx/solutions/management-pack-for-vcloud-director>.

For information on how to use the VMware vRealize Operations Tenant App, see *Using vRealize Operations Tenant App for VMware Cloud Director as a Service Provider* and *Using vRealize Operations Tenant App for VMware Cloud Director as a Tenant*.

View Use Information for a Provider Virtual Data Center

Provider virtual data centers provide compute, memory, and storage resources to its organization virtual data centers. You can monitor the use of the provider virtual data center resources, so that you can decide to add more resources.

Procedure

- 1 From the top navigation bar, select **Resources** and click **Cloud Resources**.
- 2 In the left panel, select **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Configure > Metricstab**.
- 4 For details about each parameter, click each information icon.

The Content Libraries view in the VMware Cloud Director Service Provider Admin Portal provides an interface for the integration with vRealize Orchestrator. The vRealize Orchestrator workflows are available as a catalog of services that service provider administrators can publish to tenants or other service providers and in this way extend the set of functionalities and management capabilities they offer.

This chapter includes the following topics:

- [Integrating vRealize Orchestrator with VMware Cloud Director](#)
- [Create a Service Category](#)
- [Edit a Service Category](#)
- [Import a Service](#)
- [Search for a Service](#)
- [Execute a Service](#)
- [Change a Service Category](#)
- [Unregister a Service](#)
- [Publish a Service](#)

Integrating vRealize Orchestrator with VMware Cloud Director

You integrate vRealize Orchestrator with VMware Cloud Director through the VMware Cloud Director Service Provider Admin Portal.

Integrating vRealize Orchestrator with VMware Cloud Director extends the base functionality of VMware Cloud Director by allowing service provider administrators to develop complex automation tasks through workflow orchestration and utilization of third-party plug-ins.

Through the VMware Cloud Director Service Provider Admin Portal, service provider administrators are able to view, import, and execute workflows from registered vRealize Orchestrator server instances.

In the VMware Cloud Director Service Provider Admin Portal, vRealize Orchestrator workflows can be published to service providers or tenants, allowing for quick access control and execution of both custom and built-in services.

vRealize Orchestrator has an extensive workflow library that contains pre-built tasks designed to solve specific challenges and perform common administrative tasks. Third-party plug-ins are also available at [VMware Solution Exchange](#).

Register a vRealize Orchestrator Instance with VMware Cloud Director

To leverage orchestration of workflows and automation of tasks through vRealize Orchestrator in VMware Cloud Director, you register a vRealize Orchestrator instance in the VMware Cloud Director Service Provider Admin Portal.

Prerequisites

- Deploy and configure a vRealize Orchestrator server instance. For more information, see *Installing and Configuring VMware vRealize Orchestrator* in the vRealize Orchestrator documentation.
- Configure vRealize Orchestrator to use vSphere as an authentication provider.
- Verify that VMware Cloud Director is registered with the lookup service of the same Platform Services Controller as the vCenter Single-Sign On that vRealize Orchestrator uses for authentication.

Procedure

- 1 From the top navigation bar, select **Libraries**
 - a From the left panel, select **Service Management**.
A list of registered vRealize Orchestrator server appears.
- 2 To register a new vRealize Orchestrator server, click **Add**.
The **Register vRealize Orchestrator** dialog appears.
- 3 Enter the following values.

Option	Description
Name	Name for the registered vRealize Orchestrator instance.
Description	Description for the registered vRealize Orchestrator server instance.
Hostname	The fully-qualified domain name and server port of the vRealize Orchestrator server. The default HTTPS port value is 443. Note VMware Cloud Director connects to the API interface of vRealize Orchestrator.
Username	A user account that is member of the vRealize Orchestrator administrators group.

Option	Description
Password	The password for the vRealize Orchestrator administrator account.
Trust Anchor	The vRealize Orchestrator server SSL certificate in a PEM format. Click the upload icon to find and select the .pem file.

- Click **OK** to complete the registration.

The vRealize Orchestrator server is registered with VMware Cloud Director.

Create a Service Category

You can organize services in service categories.

Procedure

- From the top navigation bar, select **Libraries**
 - From the left panel, select **Service Management**.
 - Navigate to the **Service Categories** tab.

A list of existing server categories appears.

- To create a new service category, click **Add**.

The **New Service Category** dialog appears.

- Enter the following values.

Option	Description
Name	Name of the service category.
Icon	Import the displayed icon for the service category.
Description	Short description of the service category.

Edit a Service Category

You can edit existing service categories.

Procedure

- From the top navigation bar, select **Libraries**
 - From the left panel, select **Service Management**.
 - Navigate to the **Service Categories** tab.

A list of existing server categories appears.
- Use the list bar () on the left of a selected service category and click **Edit**.

3 Edit the following values.

Option	Description
Name	Name of the service category.
Icon	Import the displayed icon for the service category.
Description	Short description of the service category.

Import a Service

You can import services from the workflow library of a vRealize Orchestrator instance that is registered with VMware Cloud Director.

Prerequisites

- Register a vRealize Orchestrator instance. See [Register a vRealize Orchestrator Instance with VMware Cloud Director](#).
- Create a service category. See [Create a Service Category](#).

Procedure

1 From the top navigation bar, select **Libraries**.

- From the left panel, select **Service Library**.

Available services display in a card view of 12 items per page, sorted by names in alphabetical order. Each card indicates that the item is a vRealize Orchestrator workflow and shows the name of the service and a tag that corresponds to the service category, in which the workflow is imported.

2 To import a new service, click the **Import** button.

3 Follow the steps of the **Import** wizard.

Option	Description
Import to target library	Select the service category, to which to import the service.
Select source	Select the vRealize Orchestrator instance, from which to import workflows.
Select workflows	Expand the hierarchical tree view to select one or multiple workflows to import.
Review	Review the details and click Done to complete the import.

The imported workflows appear in the **Service Library** card view.

Search for a Service

You can search for a service by its name or the service category it belongs to.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Service Library**.

Available services display in a card view of 12 items per page, sorted by names in alphabetical order. Each card indicates that the item is a vRealize Orchestrator workflow and shows the name of the service and a tag that corresponds to the service category, in which the workflow is imported.

- 2 In the **Search** text box on the top of the page, enter a word or a character of the name of the service or the service category you want to find.

- 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 vRealize Orchestrator workflows as imported services.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Service Library**.

Available services display in a card view of 12 items per page, sorted by names in alphabetical order. Each card indicates that the item is a vRealize Orchestrator workflow and shows the name of the service and a tag that corresponds to the service category, in which the workflow is imported.

- 2 To execute a service, in the card of the selected service, click **Execute**.

The **Execute a service** wizard appears.

- 3 Fill in the required input parameters of the service and click **Finish**.

Results

You can monitor the status of the execution in the **Recent Tasks** view. For more information, see [View Tasks](#).

Note When you start a vRealize Orchestrator workflow as a VMware Cloud Director service, VMware Cloud Director adds a few custom parameters to the workflow execution context.

Custom Property	Description
_vcd_orgName	Name of the organization, to which the user who executes the service belongs.
_vcd_orgId	ID of organization, to which the user who executes the service belongs.
_vcd_userName	Name of the user who executes the service.
_vcd_isAdmin	Has value <code>True</code> if the user who executes the service is an administrator .
_vdc_isAdmin	Deprecated. Has value <code>True</code> if the user who executes the service is an administrator .
_vdc_userName	Deprecated. Name of the user who executes the service.
_vcd_sessionToken	Authentication token you received after successful authentication to VMware Cloud Director
_vcd_apiEndpoint	VMware Cloud Director REST API endpoint

Change a Service Category

You can change the category, to which a service belongs.

Procedure

1 From the top navigation bar, select **Libraries**.

a From the left panel, select **Service Library**.

Available services display in a card view of 12 items per page, sorted by names in alphabetical order. Each card indicates that the item is a vRealize Orchestrator workflow and shows the name of the service and a tag that corresponds to the service category, in which the workflow is imported.

2 In the card of the selected service, select **Manage > Change Category**.

The **Change Category** dialog opens.

3 Select the category in which to place the service and click **Save**.

Unregister a Service

You can remove access to a service for both service providers and tenants by unregistering the service.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Service Library**.

Available services display in a card view of 12 items per page, sorted by names in alphabetical order. Each card indicates that the item is a vRealize Orchestrator workflow and shows the name of the service and a tag that corresponds to the service category, in which the workflow is imported.

- 2 In the card of the selected service, select **Manage > Unregister Workflow**.

The **Unregister Workflow** dialog opens.

- 3 To remove the service from the service library, click **Delete**.

Publish a Service

You can control service provider and tenant access to services by publishing a service.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Service Library**.

Available services display in a card view of 12 items per page, sorted by names in alphabetical order. Each card indicates that the item is a vRealize Orchestrator workflow and shows the name of the service and a tag that corresponds to the service category, in which the workflow is imported.

- 2 In the card of the selected service, select **Manage > Publish Workflow**.

The **Publish Workflow** dialog appears.

- 3 To publish to service providers, select **Publish to Service Providers** and click **Save**.

- 4 To publish to a specific tenant organization, select **Publish to Tenants** button.

- a A list with available tenant organizations appears. Select the tenant organization, to which to publish the workflow and click **Save**.

- 5 To publish to all tenant organizations, select **Publish to All Tenants** and click **Save**.

Managing Defined Entities

14

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.

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 you create an RDE type, you create a rights bundle for the type. To provide access to specific operations, you must assign rights from this bundle to other roles. Each bundle has 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.

System administrators with the **View General ACL** right can view the ACLs assigned to a specific defined entity by using the `accessControls` API. For the VMware Cloud Director API reference, see code.vmware.com.

System administrators with the **Manage General ACL** right can create, modify, and remove specific ACLs by using the `accessControls` API.

Table 14-1. Rights and ACL Entries for RDE Operations

Entity Operation	Option	Description
Read	Administrator View: TYPE right	Users with this right can see all RDEs of this type within an organization.
	View: TYPE right and ACL entry \geq View	Users with this right and a read-level ACL can view RDEs of this type.
Modify	Administrator Full Control: TYPE right	Users with this right can create, view, modify, and delete RDEs of this type in all organizations.
	Edit: TYPE right and ACL entry \geq Change	Users with this right and modify-level ACL can create, view, and modify RDEs of this type.
Delete	Administrator Full Control: TYPE right	Users with this right can create, view, modify, and delete RDEs of this type in all organizations.
	Full Control: TYPE right and ACL entry = Full Control	Users with this right and full control-level ACL can create, view, modify, and delete RDEs of this type.

You can use the VMware Cloud Director API or UI to publish the rights bundle to any organizations you want to manage the entities of this type. After publishing the rights bundle, you can assign rights from the bundle to roles within the organization.

You can use the VMware Cloud Director API to edit the ACL table.

This chapter includes the following topics:

- [Sharing Defined Entities](#)
- [Managing Custom Entities](#)

Sharing Defined Entities

You can grant access to Runtime Defined Entities (RDEs) by sharing them with other system administrators or tenants.

Sharing Defined Entities with Another User

- 1 If you want to grant access to defined entities to tenants, publish the rights bundle of the defined entity type to a tenant organization. For example, for the creation and management of Tanzu Kubernetes clusters, you must publish the **vmware:tkgcluster Entitlement** rights bundle. See [Publish or Unpublish a Rights Bundle](#).

If you want to share the defined entity with a **system administrator**, skip this step.

- 2 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.

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.

- 3 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.

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.

You can also use REST API calls to revoke access or to view who has access to the entity. See the VMware Cloud Director REST API documentation on code.vmware.com.

Sharing Administrator Rights to Defined Entities

- 1 If you want to grant access to defined entities to tenants, publish the rights bundle of the defined entity type to a tenant organization. For example, for the creation and management of Tanzu Kubernetes clusters, you must publish the **vmware:tkgcluster Entitlement** rights bundle. See [Publish or Unpublish a Rights Bundle](#).

If you want to share the defined entity with a **system administrator**, skip this step.

- 2 Assign the **Administrator View: TYPE** or **Administrator Full Control: TYPE** right from the bundle to the user roles you want to have the specific level of access to the defined entity.

For example, if you want the users with this role to view all Tanzu Kubernetes clusters within the organization, you must add the **Administrator View: Tanzu Kubernetes Guest Cluster** right to the role. If you want the users with this role to create, view, modify, and delete Tanzu Kubernetes clusters in all organizations, add the **Administrator Full Control: Tanzu Kubernetes Guest Cluster** right to the user role.

Users with the **Administrator Full Control: Tanzu Kubernetes Guest Cluster** right can grant ACL access to any VMWARE:TKGCLUSTER entity.

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.

Managing Custom Entities

The custom entity definitions in VMware Cloud Director are object types that are bound to vRealize Orchestrator object types. When a service provider publishes a custom entity definitions to either another service provider, or to one or more tenants, users VMware Cloud Director can own, manage, and change these types according to their needs. By executing services, service provider users and 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 a custom entity by its name.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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. Note 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, SSH > Host . 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 Done 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.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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.

- 3 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.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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 > Associate Action**.

A new dialog opens.

- 3 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 Service Library 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 Done 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

You can remove a vRealize Orchestrator workflow from the list of associated actions.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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.

Procedure

- 1 From the top navigation bar, select **Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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.

Procedure

1 From the top navigation bar, select **Libraries**.

a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of 12 items per page, sorted by names in alphabetical order. 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.