

vCloud Director Service Provider Admin Portal Guide

28 MAR 2019

VMware Cloud Director 9.7

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 © 2018-2020 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

1	About the vCloud Director Service Provider Admin Portal	8
	Updated Information	9
2	Getting Started with vCloud Director Service Provider Admin Portal	10
	Overview of vCloud Director Administration	10
	Log in to vCloud Director Service Provider Admin Portal	13
	View Tasks	14
	Stop a Task in Progress	14
	View Events	15
3	Managing vSphere Resources	16
	Adding vCenter Server and NSX Resources	17
	Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance	17
	Assign the NSX License Key in vCenter Server	20
	Register an NSX-T Manager Instance	20
	View the vCenter Server Instances	21
	Modify vCenter Server Settings	22
	Activate or Deactivate a vCenter Server Instance	22
	Reconnect a vCenter Server Instance	23
	Refresh a vCenter Server Instance	23
	Refresh the Storage Policies of a vCenter Server Instance	23
	Unregister a vCenter Server Instance	24
	Modify NSX Manager Settings	24
	Modify NSX-T Manager Settings	24
	Delete an NSX-T Manager Instance	25
	Multisite Resource Lists	25
4	Managing Provider Virtual Data Centers	27
	Activate or Deactivate a Provider Virtual Data Center	27
	Delete a Provider Virtual Data Center	28
	Edit the General Settings of a Provider Virtual Data Center	28
	Merge Provider Virtual Data Centers	29
	View the Organization Virtual Data Centers of a Provider Virtual Data Center	29
	View the Datastores on a Provider Virtual Data Center	30
	View the External Networks on a Provider Virtual Data Center	31
	Managing the VM Storage Policies on a Provider Virtual Data Center	31
	Add a VM Storage Policy to a Provider Virtual Data Center	32

Activate or Deactivate a VM Storage Policy on a Provider Virtual Data Center	32
Delete a VM Storage Policy from a Provider Virtual Data Center	33
Modify the Metadata for a VM Storage Policy on a Provider Virtual Data Center	33
Managing the Resource Pools on a Provider Virtual Data Center	34
Add a Resource Pool to a Provider Virtual Data Center	34
Activate or Deactivate a Resource Pool on a Provider Virtual Data Center	35
Detach a Resource Pool from a Provider Virtual Data Center	35
Modify the Metadata for a Provider Virtual Data Center	36

5 Managing Organizations 37

Understanding Leases	37
Create an Organization	38
Configure Catalogs for an Organization	38
Configure Policies for an Organization	39

6 Managing Organization Virtual Data Centers 41

Understanding Allocation Models	41
Suggested Use of the Allocation Models	43
Flex Allocation Model	44
Allocation Pool Allocation Model	45
Pay-As-You-Go Allocation Model	46
Reservation Pool Allocation Model	47
Understanding Compute Policies	48
Provider Virtual Data Center Compute Policies	48
Virtual Data Center Compute Policies	50
Create an Organization Virtual Data Center	54
Activate or Deactivate an Organization Virtual Data Center	57
Delete an Organization Virtual Data Center	57
Modify the Name and the Description of an Organization Virtual Data Center	58
Modify the Allocation Model Settings of an Organization Virtual Data Center	58
Modifying the Storage Settings of an Organization Virtual Data Center	59
Modify the VM Provisioning Settings of an Organization Virtual Data Center	59
Add a VM Storage Policy to an Organization Virtual Data Center	59
Change the Default Storage Policy on an Organization Virtual Data Center	60
Edit the Limit of a Storage Policy on an Organization Virtual Data Center	60
Modify the Metadata for a VM Storage Policy on an Organization Virtual Data Center	61
Activate or Deactivate a Storage Policy on an Organization Virtual Data Center	62
Delete a Storage Policy from an Organization Virtual Data Center	62
Edit the Network Settings of an Organization Virtual Data Center	63
Modify the Metadata for an Organization Virtual Data Center	64
View the Resource Pools of an Organization Virtual Data Center	64

Managing the Distributed Firewall on an Organization Virtual Data Center	65
Activate the Distributed Firewall on an Organization Virtual Data Center	65
Add a Distributed Firewall Rule	66
Edit a Distributed Firewall Rule	68
Custom Grouping Objects	69
Working with Security Groups	73
Working with Security Tags	76

7 Managing Edge Gateways 81

Working with Edge Clusters	81
Add an Edge Gateway	83
Configuring Edge Gateway Services	85
Managing an Edge Gateway Firewall	85
Managing Edge Gateway DHCP	89
Add a SNAT or a DNAT Rule	94
Advanced Routing Configuration	96
Load Balancing	104
Secure Access Using Virtual Private Networks	117
SSL Certificate Management	142
Custom Grouping Objects	149
View the Networks Use and IP Allocations on an Edge Gateway	153
Editing Edge Gateway Properties	153
Activate or Deactivate Distributed Routing on an Edge Gateway	154
Modify the External Networks and the Edge Gateway Settings	154
Edit the General Settings of an Edge Gateway	154
Edit the Default Gateway of an Edge Gateway	155
Edit the IP Settings of an Edge Gateway	156
Edit the Suballocated IP Pools on an Edge Gateway	156
Edit the Rate Limits on an Edge Gateway	157
Redeploy an Edge Gateway	157
Delete an Edge Gateway	157
Statistics and Logs for an Edge Gateway	158
View Statistics	158
Enable Logging	158
Enable SSH Command-Line Access to an Edge Gateway	160

8 Managing Organization Virtual Data Center Networks 161

Managing NSX-T Organization Virtual Data Center Networks	161
Add an NSX-T Organization Virtual Data Center Network	161
Edit an NSX-T Organization Virtual Data Center Network	162
Delete an NSX-T Organization Virtual Data Center Network	163

9	Managing SDDCs and SDDC Proxies	164
10	Managing System Administrators and Roles	166
	Managing Rights and Roles	166
	Predefined Roles and Their Rights	168
	New Rights in This Release	176
	Managing Rights Bundles	177
	Managing Global Tenant Roles	179
	Managing Provider Roles	182
	Managing Provider Users and Groups	184
	Managing Provider Users	184
	Managing Provider Groups	187
11	Managing System Settings	189
	Managing Identity Providers	189
	Managing LDAP Connections	189
	Configure Your System to Use a SAML Identity Provider	192
	Managing Plug-Ins	194
	Upload a Plug-in	194
	Activate or Deactivate a Plug-in	195
	Delete a Plug-in	195
	Publish or Unpublish a Plug-in from an Organization	195
	Customizing the vCloud Director Portals	196
12	Monitoring vCloud Director	198
	vCloud Director and Cost Reporting	198
	View Use Information for a Provider Virtual Data Center	198
13	Managing Services	200
	Integrating vRealize Orchestrator with vCloud Director	200
	Register a vRealize Orchestrator Instance with vCloud Director	201
	Create a Service Category	202
	Edit a Service Category	202
	Import a Service	203
	Search for a Service	203
	Execute a Service	204
	Change a Service Category	205
	Unregister a Service	205
	Publish a Service	205
14	Managing Custom Entities	207

Search for a Custom Entity	207
Edit a Custom Entity Definition	208
Add a Custom Entity Definition	208
Custom Entity Instances	209
Associate an Action to a Custom Entity	209
Dissociate an Action From a Custom Entity	210
Publish a Custom Entity	211
Delete a Custom Entity	211

About the vCloud Director Service Provider Admin Portal

1

The *vCloud 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, user, 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 vCloud Director Service Provider Admin Portal.

Related Documentation

See the *vCloud Director Administrator's Guide* for information about the features and capabilities available to an administrator using the vCloud Director Web console instead of the vCloud 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>.

Updated Information

This *vCloud Director Service Provider Admin Portal Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *vCloud Director Service Provider Admin Portal Guide*.

Revision	Description
05 APR 2019	Improved the information in chapters Understanding Allocation Models and Understanding Compute Policies .
28 MAR 2019	Initial release.

Getting Started with vCloud Director Service Provider Admin Portal

2

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

This chapter includes the following topics:

- [Overview of vCloud Director Administration](#)
- [Log in to vCloud Director Service Provider Admin Portal](#)
- [View Tasks](#)
- [Stop a Task in Progress](#)
- [View Events](#)

Overview of vCloud Director Administration

With VMware vCloud 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 *vCloud Director Administrator's 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

vCloud 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. vCloud Director also uses vSphere distributed switches, vSphere port groups, and NSX Data Center for vSphere to support virtual machine networking.

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

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

Starting with version 9.7, vCloud 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 vCloud 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. In addition, vCloud Director 9.7 introduces the Software-Defined Data Center (SDDC) and the SDDC proxies as cloud resources that provide access to the underlying vSphere environment from vCloud Director.

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

SDDCs and SDDC Proxies

vCloud Director 9.7 introduces the SDDC as a cloud resource that encapsulates an entire vCenter Server installation. An SDDC includes one or more SDDC proxies that are access points to different components of the underlying vSphere environment. The provider can create and enable an SDDCs and proxies. The provider can publish an SDDC and its proxies to tenants.

To create and manage SDDCs and proxies, you must use the vCloud OpenAPI. See *Getting Started with vCloud 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.

vCloud Director Networking

vCloud 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, vCloud 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 vCloud 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.

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

Starting with vCloud 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.

vCloud 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 *vCloud 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. vCloud 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 vCloud 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

vCloud 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 *vCloud 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 vCloud Director Service Provider Admin Portal

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

Prerequisites

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

Procedure

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

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

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


View Tasks

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

The tasks view is a good way to view the status of tasks in your service provider admin portal at a glance. The view shows when the tasks were executed, and whether they were successful. This tool can be a good first step for troubleshooting any issues in your environment.

The blue and red info tips over the Tasks icon show the number of run and failed tasks, respectively.

Procedure

- ◆ From the upper-right menu, select the Tasks icon (.

Results

A lists of recent tasks displays, along with the time the task was executed and the status of the task.

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 From the main menu () , select **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>vdUpdateVapp</i> .
Status	Status of the event, such as Succeeded or Failed.
Service namespace	Service name, such as <i>com.vmware.vcloud</i> .
Organization	Name of the organization.
Owner	User who triggered the event.
Time of occurrence	Date and time when the event occurred.

Managing vSphere Resources

3

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

vCloud Director uses one or more vCenter Server environments to back its virtual data centers. Starting with version 9.7, vCloud 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 vCloud Director with their vCloud Director accounts.

Before you can use a vCenter Server instance in vCloud 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 service provider, also called provider scoped. For information about creating a provider virtual data center, see the *vCloud Director Administrator's Guide*.

When you create an SDDC that encapsulates an attached vCenter Server instance, this vCenter Server instance appears as published to tenants, also called tenant scoped. For information about creating an SDDC, see [Chapter 9 Managing SDDCs and SDDC Proxies](#).

Note By default, with an attached vCenter Server instance, you can create either a provider VDC or an SDDC. If you created a provider VDC backed by an vCenter Server instance, you cannot use this vCenter Server instance to create an SDDC, and the reverse. You can use the vCloud API to modify the system settings of your vCloud Director installation so that a vCenter Server instance can back both a provider VDC and an SDDC.

This chapter includes the following topics:

- [Adding vCenter Server and NSX 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](#)
- [Multisite Resource Lists](#)

Adding vCenter Server and NSX Resources

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

For information about vCloud 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 attach a vCenter Server instance so that its resources become available for use in vCloud Director. You can attach a vCenter Server instance together with its associated NSX Manager instance, or you can attach a vCenter Server instance alone.

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

If you want vCloud 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 vCloud 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 vCloud Director environment.

Prerequisites

- If you configured vCloud Director to verify vCenter and vSphere SSO certificates, verify that you uploaded the vCenter Server certificates to vCloud Director. For information about the general system settings, see the *vCloud Director Administrator's Guide*.

- If you configured vCloud Director to verify NSX Manager certificates, verify that you uploaded the NSX Manager certificates to vCloud Director. For information about the general system settings, see the *vCloud Director Administrator's Guide*.

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 vCloud 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 main menu () , select **vSphere Resources**.
- 2 In the left pane, click **vCenters** and click **Add**.
- 3 If you have a multisite vCloud 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 vCloud 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 vCloud 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, https://example.vmware.com/vsphere-client .

- 9 Click **Next**.

- 10 (Optional) Skip adding the NSX Manager instance that is associated with the vCenter Server instance, and finish the registration.

If you want vCloud 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.

- a On the **NSX-V Manager Settings** page, turn off the **Configure Settings** toggle, and click **Next**.
- b On the **Ready to Complete** page, review the registration details and click **Finish**.

(Optional) Add the Associated NSX Manager Instance

If you want vCloud 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 Settings** 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
Resource Pool Path	<p>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.</p> <p>As an alternative, you can enter the Managed Object Reference ID of the resource pool. For example, resgroup-1476.</p>
Datastore Name	The name of the datastore to host the appliance files. For example, shared-disk-1 .

Option	Description
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>vCloud Director Tenant Portal Guide</i> .

- 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](#).
- For information about creating a provider virtual data center, see the *vCloud Director Administrator's Guide*.

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 vCloud Director networking.

Prerequisites

This operation is restricted to system administrators.


Procedure

- From a vSphere Client that is connected to the vCenter Server system, select **Home > Licensing**.
- For the report view, select **Asset**.
- Right-click the NSX Manager asset and select **Change license key**.
- Select **Assign a new license key** and click **Enter Key**.
- 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 vCloud Director. You can use this license key in multiple vCenter Server instances.
- Click **OK**.

Register an NSX-T Manager Instance

You can register an NSX-T Manager instance with vCloud Director, so that vCloud 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 main menu () , select **vSphere Resources**.
- 2 In the left pane, click **NSX-T Managers** and click **Add**.
- 3 If you have a multisite vCloud 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 vCloud 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 *vCloud API Programming Guide for Service Providers* at <https://code.vmware.com>.

View the vCenter Server Instances

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

Procedure

- 1 From the main menu () , select **vSphere Resources**.
- 2 In the left panel, click **vCenters**.

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 vCloud Director.
State	Activated or deactivated. See Activate or Deactivate a vCenter Server Instance .
Connection	Connected or not to vCloud Director. See Reconnect a vCenter Server Instance .
VC Host	FQDN of the vCenter Server instance.
Version	The vCenter Server version.


	Description
Service Provider	Published or not for use by virtual data centers.
Tenant	Published or not for use as a software-defined data center (SDDC).
Site	The vCloud Director FQDN for the site to which the vCenter Server instance belongs.

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 vCloud Director, 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 main menu () , select **vSphere Resources**.
- 2 In the left pane, click **vCenters** and click the name of the vCenter Server instance that you want to modify.
- 3 In the upper right corner of the **vCenter Info** section, click **Edit**.
- 4 Edit the vCenter Server settings, and 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 vCloud Director, you must activate the vCenter Server instance.

Procedure


- 1 From the main menu () , select **vSphere Resources**.
- 2 In the left panel, click **vCenters**.
- 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 main menu () , select **vSphere Resources**.
- 2 In the left panel, click **vCenters**.
- 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 vCloud Director database about the underlying vCenter Server resources, you must refresh the vCenter Server instance.


Procedure

- 1 From the main menu () , select **vSphere Resources**.
- 2 In the left panel, click **vCenters**.
- 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 vCloud 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 main menu () , select **vSphere Resources**.
- 2 In the left panel, click **vCenters**.
- 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 vCloud 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 main menu () , select **vSphere Resources**.
- 2 In the left panel, click **vCenters**.
- 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 vCloud 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 main menu () , select **vSphere 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 Edit the vCenter Server settings, and click **Save**.

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 vCloud 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

- 1 From the main menu () , select **vSphere Resources**.
- 2 In the left pane, click **NSX-T Managers** and click the name of the NSX-T Manager instance that you want to modify.
- 3 In the upper right corner of the **General** tab, click **Edit**.
- 4 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 vCloud 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 main menu () , select **vSphere 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**.

Multisite Resource Lists

If you are working with vCloud 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, vCloud Director introduces multisite resource lists.

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 organization VDCs, organization VDC templates, provider VDCs, cloud cells, edge gateways, external networks, and network pools.

Note Multisite organizations lists are not supported.

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 vCloud Director Web Console or the vCloud API. For information about using vCloud Director Web Console, see the *vCloud Director Administrator's Guide*. For information about using the vCloud API, see the *vCloud API Programming Guide for Service Providers*.

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](#)
- [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 main menu () , select **Cloud Resources**.

- 2 In the left panel, click **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 vCloud 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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 main menu () , select **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) 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.

- 6 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 site.
- The target provider virtual data centers contain only elastic organization virtual data centers.


Procedure

- 1 From the main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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 vCloud 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Datastore** tab.

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

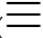
Title	Description
Name	The name of the datastore
State	Enabled or disabled
Type	The type of file system that the datastore uses, either Virtual Machine File System (VMFS) or Network File System (NFS)
Used	The datastore space occupied by virtual machine files, including log files, snapshots, and virtual 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.

Title	Description
Requested Storage	<p>Provisioned storage in use only by vCloud Director objects on the datastore, including:</p> <ul style="list-style-type: none"> ■ Virtual machines provisioned in vCloud 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	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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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.

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

Add a VM Storage Policy to a Provider Virtual Data Center

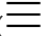
You can add a VM storage policy to a provider virtual data center, after which you can configure organization virtual data centers backed by this provider virtual data center to support the added storage policy.

Important vCloud Director does not support VM storage policies for host-based data services such as encryption and Storage I/O Control.

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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 On the **Storage Policies** tab, click **Add**.
- 4 Select one or more storage policies that you want to add, and click **Add**.

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

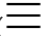
What to do next

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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Storage Policies** tab.
- 4 Click the radio button next to the target VM storage policy, and click **Enable** or **Disable**.
- 5 To confirm, click **OK**.

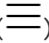
Delete a VM Storage Policy from a Provider Virtual Data Center

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

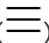
- 1 From the main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 Click the **Storage Policies** tab.
- 4 Click the radio button next to the target VM storage policy, and click **Remove**.
- 5 To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

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.

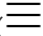
When compute resources are backed by multiple resource pools, they can expand to accommodate more virtual machines.

You can add resource pools backed by vSphere clusters that are optimally configured for hosting NSX edges that have VLAN uplinks. vCloud 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Provider VDCs**, and click the name of the target provider virtual data center.
- 3 On the **Resource Pools** tab, click **Add**.
- 4 Select one or more resource pools to add, and click **Add**.

Results

vCloud 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.

vCloud 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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](#).
- Migrate any virtual machines from that resource pool to an activated resource pool. For information about migrating virtual machines between resource pools on a provider virtual data center, see the *vCloud Director Administrator's Guide*.
- 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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 vCloud Director Service Provider Admin Portal allows you to create, configure, and manage vCloud Director organizations.

Use vCloud 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](#)
- [Configure Catalogs for an Organization](#)
- [Configure Policies for 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, vCloud 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, vCloud 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 vCloud Director Service Provider Admin Portal.

Procedure

- 1 From the main menu (), select **Cloud Resources**
 - a From the left panel, select **Organizations**.

The list of existing organizations displays in a grid view.

- 2 To create a new organization, click the **+Add** button.

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.


Configure Catalogs for an Organization

You can configure how an organization shares its service catalogs.

Procedure

- 1 From the main menu (), select **Cloud Resources**
 - a From the left panel, select **Organizations**.

The list of existing organizations displays in a grid view.

- 2 Use the list bar () on the left of each item to display the actions you can take for each organization.

- 3 Click **Catalogs**.

The organization's **Catalog Settings** dialog opens.

4 Configure the following sharing and publishing options.

Option	Description
Sharing	Allows organization administrators to share this organization's catalogs with other organizations in this instance of vCloud 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 vCloud Director.
Allow subscribing to external catalogs	Allows organization administrators to subscribe to catalogs outside of this instance of vCloud 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 main menu (), select **Cloud Resources**
 - a From the left panel, select **Organizations**.

The list of existing organizations displays in a grid view.

- 2 Use the list bar () on the left of each item to display the actions you can take for each organization.
- 3 Click **Policies** to edit the leases, quotas, resource limits, and password policies for the organization.
- 4 Configure vApp leases with the following settings.

Option	Description
Maximum runtime lease	How long vApps can run before they are automatically stopped.
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.

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 Compute Policies](#)
- [Create an Organization Virtual Data Center](#)
- [Activate or Deactivate an Organization Virtual Data Center](#)
- [Delete an Organization Virtual Data Center](#)
- [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](#)
- [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 VDC compute 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 VDC Compute policy of the VM.	Unlimited	Unlimited	vCPU speed times the number of vCPUs	Custom
VM CPU Reservation	Based on the VDC Compute 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 VDC Compute policy of the VM.	Unlimited	Unlimited	vRAM	Custom
VM RAM Reservation	Based on the VDC Compute policy of the VM.	0	Equals vRAM times RAM guarantee plus RAM overhead.	Equals vRAM times RAM guarantee plus RAM overhead.	Custom

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, vCloud 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.

Allocation model	Suggested use
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 vCloud 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 vCloud 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 VDC compute 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.

If you create a non-flex organization VDC in vCloud Director 9.7, you can reconfigure the organization VDC to use the flex allocation model. If an organization VDC is created by using vCloud Director version earlier than 9.7, you cannot reconfigure the organization data centers to use the flex allocation model.

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

- Activate or deactivate the elastic pool feature.
- Include or exclude memory overhead.
- Specify a default VDC compute policy for the organization VDC.
- Memory and CPU allocation and guarantee
- Network quota
- Storage profile

As a **vCloud 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 vCloud 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 organization VDC compute policies without any constraints that other allocation models have. In the flex allocation model, the VM compute resource allocation depends on the organization VDC compute policies. If you do not define a VDC compute 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 VDC compute policies, a single organization VDC can accommodate VMs that use configuration that is common for all other allocation models. For more information, see [Understanding Compute Policies](#).

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

Allocation Pool Allocation Model

With the allocation pool allocation model, 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. This percentage is known as the percentage guarantee factor, and it allows you to overcommit resources.

Starting with vCloud Director 5.1.2, system administrators can configure allocation-pool organization VDCs to be elastic or non-elastic. Elasticity is a global setting that affects all allocation-pool organization VDCs. For information about modifying general system settings, see the *vCloud Director Administrator's Guide*.

By default, allocation-pool organization VDCs have an elastic allocation pool enabled. Systems upgraded from vCloud 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 vCloud 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 vCloud 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 *vCloud API Programming Guide for Service Providers*.

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, vCloud 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 *vCloud API Programming Guide for Service Providers*.

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 Compute Policies

Starting with vCloud Director 9.7, you can control the resource allocation and the virtual machine (VM) placement by using compute policies. Based on the scope and the function, there are two types of compute policies - provider virtual data center (VDC) compute policies and VDC compute policies.

Provider VDC compute policy

A provider VDC compute policy defines VM-host affinity rules that directly impact the placement of tenant workloads. Tenant users have no visibility over the provider VDC compute policies.

The scope of provider VDC compute policies is at the provider VDC level.

VDC compute policy

A VDC compute policy controls the compute characteristics of a VM at the organization VDC level. Because tenant users have no visibility over the provider VDC compute policies, to expose the VM-host affinity rules for tenant use, you refer the provider VDC compute policy inside the VDC compute policy.

Provider Virtual Data Center Compute Policies

By using provider virtual data center (VDC) compute policies, vCloud Director **system administrators** can expose virtual machine (VM) groups and logical VM groups to tenants.

Provider VDC compute policies might contain a collection of the following:

- VM groups that contain similar VMs. Each VM group belongs to a different cluster.
- Logical VM groups that are suited for diverse functionalities.
- Both VM groups and logical VM groups.

Provider VDC Compute Policies and Logical VM Groups

System administrator can expose vSphere Distributed Resource Scheduler (DRS) VM-host affinity rules to tenants by using VM groups and logical VM groups. DRS VM-host affinity rules are exposed at the provider level in vCloud Director as VM groups. VM-host affinity rules are bound to a specific cluster. Because elastic provider VDCs can span across multiple vSphere clusters, logical VM groups provide the abstraction of DRS VM-host affinity rules that works across multiple clusters by grouping cluster bound VM groups that are logically equivalent. To manage logical VM groups, you use vCloud OpenAPI. For information about vCloud OpenAPI, see *Getting Started with vCloud OpenAPI* at <https://code.vmware.com>.

To expose VM-host affinity rules, you add VM groups and logical VM groups to a provider VDC compute policy and create a reference between the provider VDC compute policy and a VDC compute policy.

In the provider VDC compute policy context, logical VM groups have an **AND** relationship between one another.

With provider VDC compute policies and logical VM groups, **vCloud Director system administrators** can expose multiple VM groups to tenant users within an organization VDC. For example, consider an environment that contains two clusters: *cluster1* and *cluster2*. In *cluster1* resides the host *SQL_host_1*, while in *cluster2* reside the hosts *SQL_fast_host* and *Fast_host*.

- 1 In *cluster1*, you create *SQL_host_group1* and *VM_group1*.

You create a positive affinity between *VM_group1* and *SQL_host_group1*.

- 2 In *cluster2*, you create four groups.

- You create *SQL_host_group2* and *VM_group2*

You create a positive affinity between *VM_group2* and *SQL_host_group2*.

- You create *fast_host_group* and *VM_group3*.

You create a positive affinity between *VM_group3* and *fast_host_group*.

You create the *PVDC_compute_policy1* that consists of *logical_VM_group1* and *logical_VM_group2*. The *logical_VM_group1* comprises *VM_group1* and *VM_group2*. The *logical_VM_group2* comprises *VM_group3*.

You create and publish the *SQL_and_fast* VDC compute policy to an organization VDC, and add a reference to *PVDC_compute_policy1*. When you create a reference between the *SQL_and_fast* VDC compute policy and the *PVDC_compute_policy1*, you expose logical VM groups and VM groups information to tenant users within the organization VDC. As a result, when a tenant applies the *SQL_and_fast* VDC compute policy to a VM, the placement engine adds the VM to the *SQL_fast_host* within *cluster2*.

The workflow is the following.

- 1 A **vCenter Server administrator** creates host groups by using the vSphere Client.

For information, see the *Create a Host DRS Group (MSCS)* topic in the *VMware vSphere ESXi and vCenter Server Documentation*.

- 2 A **vCenter Server administrator** or a **vCloud Director system administrator** creates VM groups.

For information, see the *Create or Update a VM Group* topic in the *vCloud Director Administrator's Guide*.

- 3 A **vCloud Director system administrator** creates the appropriate affinity rules between VM groups and host groups.

For information, see *Managing VM-Host Affinity Rules* topic in the *vCloud Director Administrator's Guide*.

- 4 A **vCloud Director system administrator** groups logically equivalent VM groups into logical VM groups by using the vCloud OpenAPI.

- 5 A **vCloud Director system administrator** creates a provider VDC compute policy and adds the logical VM groups by using the vCloud OpenAPI.

- 6 A **vCloud Director system administrator** creates a VDC compute policy that refers to the provider VDC compute policy, and publishes the VDC compute policy to an organization VDC by using the vCloud OpenAPI.

When a tenant creates a VM in the organization VDC and selects the VDC compute policy, vCloud Director adds the VM to the VM group that is referenced in the VDC compute policy. As a result, vCloud Director creates the VM on the appropriate host.

Provider VDC Compute Policies and VM Groups

A provider VDC compute policy can have zero or one VM group from each cluster. For example, the provider VDC compute 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 provider VDC compute policy to a VM, the placement engine adds this VM to the corresponding VM group of the cluster on which it resides. For example, if you select to deploy a VM on cluster *oracle_cluster1* and assign the provider VDC compute policy *oracle_license* to this VM, the placement engine adds the VM to VM group *oracle_license1*.

The workflow is the following.

- 1 A **system administrator** creates one or more provider VDC compute policies by using the vCloud OpenAPI.
- 2 A **system administrator** creates one or more VDC compute policies by using the vCloud OpenAPI.

A VDC compute policy can be associated with zero or one provider VDC compute policy. VDC compute policies are unique by name and by provider VDC compute policy.

- 3 A **system administrator** publishes the VDC compute policy to one or more organization VDCs by using the vCloud OpenAPI.

Tenants can see only the VDC compute policies that are published to their organization VDCs. Provider VDC compute policies are not available at a tenant level.

- 4 Tenants can use the vCloud API or the vCloud Director Tenant Portal to assign an organization VDC compute policy to a VM when creating or updating a VM.

Initially, the system does not contain any provider VDC compute policies, and each organization VDC contains only a default compute policy, which is not associated with a provider VDC compute policy.

To create and manage provider and global VDC compute policies, you must use the vCloud OpenAPI. See *Getting Started with vCloud OpenAPI* at <https://code.vmware.com>.

Virtual Data Center Compute Policies

Virtual data center (VDC) compute policies control the physical compute resource allocation for tenant workloads. To allocate physical resources based on specific workload requirements, tenant users can select between a default and custom VDC compute policies.

A VDC compute policy groups attributes that define the compute resource allocation for virtual machines within an organization VDC. The compute resource allocation includes CPU and memory allocation, reservations, limits, and shares.

vCloud Director **system administrators** create and manage compute policies at a global level and can publish individual compute policies to one or more organization VDCs. When you publish a VDC compute 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, **tenant administrators** can assign the available VDC compute policies to virtual machines. **Tenant administrators** and users in the organization VDC cannot look into the specific configuration of a VDC compute policy.

With VDC compute policies, cloud providers can define named CPU and memory consumption profiles that tenants can associate with virtual machines within an organization VDC. Using VDC Compute policies is 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. With VDC compute policies, cloud providers can also limit or constrain CPU and memory consumption of virtual machines in an organization VDC.

With VDC compute policies, vCloud 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

Attributes of Virtual Data Center Compute Policies

When you create a virtual data center (VDC) compute policy, you can specify a subset of all available attributes. The only mandatory attribute is the VDC compute policy name.

The following table lists all attributes that you can define within a VDC Compute policy.

Table 6-1. VDC Compute Policy Attributes

VDC Compute Policy Attribute	API Parameter	Description
Name	name	Mandatory parameter that is used as an identifier for the VDC compute policy.
Description	description	Represents a short description of the VDC compute policy.
vCPU Speed	cpuSpeed	Defines the vCPU speed of a virtual machine (VM) in MHz.
Memory	memory	Defines the memory configured for a VM in MB. When a tenant assigns the VDC compute policy to a VM, the VM receives the amount of memory defined by this attribute.
Number of vCPUs	cpuCount	Defines the number of vCPUs configured for a VM. When a tenant assigns the VDC compute policy to a VM, the VM receives the number of vCPUs defined by this attribute.

Table 6-1. VDC Compute Policy Attributes (continued)

VDC Compute Policy Attribute	API Parameter	Description
Cores per Socket	coresPerSocket	<p>The number of cores per socket for a VM.</p> <p>The number of vCPUs that is defined in the VDC compute policy must be divisible by the number of cores per socket.</p> <p>If the number of vCPUs is not divisible by the number of cores per socket, the number of cores per socket becomes invalid.</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 1.</p> <p>Value of 0 memory reservation guarantee defines no memory guarantee. Value of one defines 100% memory reserved.</p>
CPU Reservation Guarantee	cpuReservationGuarantee	<p>Defines how much of the CPU resources of a VM are reserved.</p> <p>The allocated CPU for a VM equals the number of vCPUs times the vCPU speed in MHz.</p> <p>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.</p>
CPU Limit	cpuLimit	<p>Defines the CPU limit in MHz for a VM.</p> <p>Value of minus one (-1) defines no CPU limit.</p> <p>If not defined in the VDC compute policy, CPU limit is equal to the allocated CPU for the VM.</p>
Memory Limit	memoryLimit	<p>Defines the memory limit in MB for a VM.</p> <p>Value of minus one (-1) defines no memory limit.</p> <p>If not defined in the VDC compute policy, memory limit is equal to the allocated memory for the VM.</p>
CPU Shares	cpuShares	<p>Defines the number of CPU shares for a VM.</p> <p>If not defined in the VDC compute policy, normal shares are applied to the VM.</p>
Memory Shares	memoryShares	<p>Defines the number of memory shares for a VM.</p> <p>If not defined in the VDC compute policy, normal shares are applied to the VM.</p>
Extra Configurations	extraConfigs	Represents a mapping between a key and value pairs that are applied as extra configuration values on a VM.
Provider VDC Compute Policy	pvdcComputePolicy	Defines the reference of the VDC compute policy to a provider VDC compute policy.

Working with Virtual Data Center Compute Policies

vCloud Director generates a default compute policy for all virtual data centers (VDCs). The default VDC compute policy contains only a name and description, and all remaining VDC compute policy attributes are empty.

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

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

If you publish multiple VDC compute 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 available VDC compute policy operations for cloud providers are the following:

- Create a VDC compute policy.
- Publish a VDC compute policy to one or more organization VDC.
- Unpublish a VDC compute policy from an organization VDC.
- Delete a VDC compute policy.

Users that have the **ORG_VDC_MANAGE_COMPUTE_POLICIES** right can create, update, and publish VDC compute policies. To create VDC compute policies, you use the vCloud API.

The following table lists the available VDC compute policy operations for tenant users.

Table 6-2. VDC Compute Policy Operations for Tenant Users

Operation	Description
Assign a VDC compute policy to a VM during a VM creation.	Tenant users that are authorized to create VMs in an organization VDC can optionally assign VDC compute policies to VMs. As a result, the parameters defined in the VDC compute policy control the CPU and memory consumption of the VM. Assigning a VDC compute policy is not a requirement for tenants during a VM creation. If a tenant does not explicitly select a VDC compute policy to assign to a VM, the default VDC policy is applied to the VM. Tenant users can assign a VDC compute policy to a VM during a VM creation using the vCloud Director Tenant Portal.
Assign a VDC compute policy to an existing VM.	Tenant users that are authorized to manage VMs in an organization VDC can update the association between a VM and a VDC compute policy. As a result, the system reconfigures the VM to consume compute resources as specified in the new VDC compute policy. Tenant users can assign a VDC compute policy to existing VM using the vCloud Director Tenant Portal.

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

- 1 A **system administrator** creates three VDC compute 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 VDC compute policies to an organization VDC.

Publishing a VDC compute policy to an organization VDC makes the policy available to tenant users in the organization VDC.

- 3 Optionally define one of the VDC compute policies as a default VDC policy for the organization VDC.

If you define a default policy for the organization VDC, and if the tenant users do not specify another policy during the creation of a VM, the default policy is applied to the VM.

To view and modify VDC compute policies, you must use the vCloud API.


Create an Organization Virtual Data Center

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

Prerequisites

Create a provider VDC. See the *vCloud Director Administrator's Guide*.

Procedure

- 1 From the main menu () , select **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.
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

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 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
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.

vCloud 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**.

vCloud 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.

- 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 main menu () , select **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 main menu () , select **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**.

Modify the Name and the Description of an Organization Virtual Data Center

As your vCloud 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 main menu () , select **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 main menu () , select **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.

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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 On the **Storage** tab, in the upper right corner, 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Storage** tab, 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.

vCloud 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Storage** tab.
- 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Storage** tab.
- 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Storage** tab.
- 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Storage** tab.
- 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organization VDCs**, and click the name of the target organization virtual data center.
- 3 Click the **Storage** tab.
- 4 Click the radio button next to the name of the target storage policy, and click **Remove**.
- 5 To confirm, click **Remove**.

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 main menu () , select **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.
 - 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 vCloud 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 *vCloud 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 main menu () , select **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 main menu () , select **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.

vCloud 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 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.

vCloud Director supports distributed firewall services on organization virtual data centers that are backed by NSX Data Center for vSphere.

Procedure

- 1 From the main menu () , select **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 main menu () , select **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: a Select the service protocol. b 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: a Select one or more objects and add them to the filter. b 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.

Note 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.


- 14 Click **Save Changes**.

Edit a Distributed Firewall Rule


In a vCloud 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 main menu () , select **Cloud Resources**.
- In the left panel, click **Organization VDCs**.
- 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 vCloud 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.

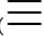
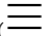
Create an IP Set for Use in Firewall Rules and DHCP Relay Configuration

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 main menu () , 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 main menu () , 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.

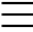
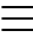
Create a MAC Set for Use in Firewall Rules

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"> a From the main menu () , 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 main menu () , 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 **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.

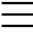
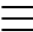
View Services Available for Firewall Rules

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"> From the main menu () , 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 main menu () , 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 **Services** tab.

Results

The available services are displayed on the screen.

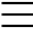
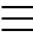
View Service Groups Available for Firewall Rules

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 main menu () , 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 main menu () , 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 main menu () , select **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 Click **Keep** to preserve your changes.

The operation can take a minute to complete.


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 main menu () , select **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 main menu () , select **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 main menu () , select **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 main menu () , select **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 main menu () , select **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 main menu () , select **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 main menu () , select **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 Edge Gateways

7

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

Edge gateways require NSX Data Center for vSphere. For information, see the *NSX Administration Guide*.

Starting with vCloud 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 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, vCloud 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 Edge Clusters](#)
- [Add an Edge Gateway](#)
- [Configuring 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 Edge Clusters

To isolate the compute workloads from the networking workloads, vCloud Director 9.7 introduces 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.

With version 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 vCloud 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 vCloud OpenAPI `EdgeClusters` methods. See *Getting Started with vCloud 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 vCloud OpenAPI `VdcNetworkProfile` methods. See *Getting Started with vCloud 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 Edge Gateway

An 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.

Starting with vCloud Director 9.7, edge gateways are deployed on edge clusters that you previously created and assigned to the organization virtual data center.

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, distribute firewalls, and static routing.

Prerequisites

- For information about the system requirements for deploying an 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 Edge Clusters](#).

Procedure

- 1 From the main menu () , select **Cloud Resources**.
- 2 In the left pane, click **Edge Gateways**, and click **New**.
- 3 Click the radio button next to the name of the 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**.

Option	Description
Compact	Requires less memory and fewer compute resources.
Large	Provides increased capacity and performance than with the Compact option. 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 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.
- 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 specify 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"> 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. 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. Click Save. The system combines overlapping IP ranges. 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. vCloud 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 vCloud 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 Edge Gateway Services

You can configure services such as DHCP, firewall, network address translation (NAT), and VPN on an Edge Gateway.

Managing an 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 deactivate or activate an edge gateway firewall.


Add an Edge Gateway Firewall Rule

You use the edge gateway Firewall screen 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 edge 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 main menu () , select **Cloud Resources**.
 - 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 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 Edge Gateway Firewall Rule](#).


Procedure

- 1 Open Edge Gateway Services.
 - a From the main menu () , select **Cloud Resources**.
 - 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 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 main menu () , select **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 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 DHCP-activated edge gateway 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 advanced 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 main menu () , select **Cloud Resources**.
 - 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

vCloud 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.

- From the main menu () , select **Cloud Resources**.
- In the left panel, click **Edge Gateways**.
- 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 Edge Gateways

The DHCP relay capability provided by NSX in your vCloud Director environment enables you to leverage your existing DHCP infrastructure from within your vCloud 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 in synch 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 Edge Gateway

The NSX software in your vCloud Director environment provides the capability for the edge gateway to relay DHCP messages to DHCP servers external to your vCloud 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 main menu () , select **Cloud Resources**.
 - 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 vCloud 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 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 main menu () , select **Cloud Resources**.
 - 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. 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.
Description	(Optional) Type a description that helps identify what this rule is doing.
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. 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.

Option	Description
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.
Description	(Optional) Type a description that helps identify what this rule is doing.
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 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 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 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 main menu () , select **Cloud Resources**.
- 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 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 edge gateway interfaces. Otherwise, configuration of that static route fails.

Procedure

- 1 Open Edge Gateway Services.
 - a From the main menu () , select **Cloud Resources**.
 - 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.

Option	Description
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 Edge Gateway Firewall Rule](#).

Configure OSPF

You can configure the Open Shortest Path First (OSPF) routing protocol for the dynamic routing capabilities of an edge gateway. A common application of OSPF on an edge gateway in a vCloud Director environment is to exchange routing information between edge gateways in vCloud 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 Edge Gateway](#).

Procedure

- Open Edge Gateway Services.
 - From the main menu () , select **Cloud Resources**.
 - 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 **Routing > OSPF**.
- 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	<p>Select Normal or NSSA.</p> <p>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.</p>
Area Authentication	<p>Select the type of authentication for OSPF to perform at the area level.</p> <p>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.</p> <p>Choices are:</p> <ul style="list-style-type: none"> ■ None <p>No authentication is required.</p> ■ Password <p>With this choice, the password you specify in the Area Authentication Value field is included in the transmitted packet.</p> ■ MD5 <p>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.</p>

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 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 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 main menu () , select **Cloud Resources**.
 - 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.

Option	Description
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.
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 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 edge gateway.

Procedure

- 1 Open Edge Gateway Services.
 - a From the main menu () , select **Cloud Resources**.
 - 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 utilization, maximizing throughput, minimizing response time, and avoiding overload.

About Load Balancing

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 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. See the *vCloud Director User's Guide*.

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 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 backend servers.

Server Pool Member

Represents the backend 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 may use an already 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 main menu () , select **Cloud Resources**.
 - 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>
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**.

The save operation can take a minute to complete.

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 main menu () , select **Cloud Resources**.
- 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) Type 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>
Cookie Name	(HTTP and HTTPS) If you specified Cookie as the persistence mechanism, type 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	(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.
Enable Pool Side SSL	(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.

- 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 Click **Keep** to preserve your changes.

The operation can take a minute to complete.



What to do next

Add service monitors for the load balancer to define health checks for different types of network traffic. See [Create a Service Monitor](#).

Create a Service Monitor

You create a service monitor to define health check parameters for a particular type of network traffic. When you associate a service monitor with a pool, the pool members are monitored according to the service monitor parameters.

Procedure

- 1 Open Edge Gateway Services.
 - a From the main menu () , select **Cloud Resources**.
 - 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 > Service Monitoring**.
- 3 Click the **Create** () button.
- 4 Enter a name for the service monitor.
- 5 (Optional) Configure the following options for the service monitor:

Option	Description
Interval	Enter the interval at which a server is to be monitored using the specified 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.

Option	Description
Type	<p>Select the way in which you want to send the health check request to the server—HTTP, HTTPS, TCP, ICMP, or UDP.</p> <p>Depending on the type selected, the remaining options in the New Service Monitor dialog are activated or deactivated.</p>
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	<p>(HTTP and HTTPS) Enter the URL to be used in the server status request.</p> <p>Note When you select the POST method, you must specify a value for Send.</p>
Send	(HTTP, HTTPS, UDP) Enter the data to be sent.
Receive	<p>(HTTP, HTTPS, and UDP) Enter the string to be matched in the response content.</p> <p>Note When Expected is not matched, the monitor does not try to match the Receive content.</p>
Extension	<p>(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:</p> <pre><extension>delay=2 critical=3 escape</extension></pre>

6 Click **Keep** to preserve your changes.

The operation can take a minute to complete.

Example: Extensions Supported for Each Protocol

Table 7-1. Extensions for HTTP/HTTPS Protocols

Monitor Extension	Description
no-body	<p>Does not wait for a document body and stops reading after the HTTP/HTTPS header.</p> <p>Note An HTTP GET or HTTP POST is still sent; not a HEAD method.</p>
max-age= <i>SECONDS</i>	Warns when a document is more than SECONDS 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 STRING.
eregi= <i>STRING</i>	Searches the page for case-insensitive regex STRING.

Table 7-1. Extensions for HTTP/HTTPS Protocols (continued)

Monitor Extension	Description
invert-regex	Returns CRITICAL when found and OK when not found.
proxy-authorization= <i>AUTH_PAIR</i>	Specifies the username:password on proxy servers with basic authentication.
useragent= <i>STRING</i>	Sends the string in the HTTP header as User Agent.
header= <i>STRING</i>	Sends any other tags in the HTTP header. Use multiple times for additional headers.
onredirect=ok warning critical follow sticky stickyport	Indicates how to handle redirected pages. <i>sticky</i> is like <i>follow</i> but stick to the specified IP address. <i>stickyport</i> ensures the port stays the same.
pagesize= <i>INTEGER:INTEGER</i>	Specifies the minimum and maximum page sizes required in bytes.
warning=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= <i>AUTH_PAIR</i>	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 \n, \r, \t, or \ in a send or quit string. Must come before a send or quit option. By default, nothing is added to send and \r\n is added to the end of quit.
all	Specifies all expect strings need to occur in a server response. By default, any 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 ok, warn, or crit. By default, uses state crit.
mismatch=ok warn crit	Accepts expected string mismatches with states ok, warn, or crit. By default, uses state warn.

Table 7-3. Extensions for TCP Protocol (continued)

Monitor Extension	Description
jail	Hides output from the TCP socket.
maxbytes= <i>INTEGER</i>	Closes the connection when more than the specified number of bytes are received.
delay= <i>INTEGER</i>	Waits the specified number of seconds between sending the string and polling for a response.
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 main menu () , select **Cloud Resources**.
 - 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.
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 host 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 Click **Keep** to preserve your changes.

The operation can take a minute to complete.


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 main menu () , select **Cloud Resources**.
 - 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/configuration-1.5.html>.

- 6 Click **Keep** to preserve your changes.

The operation can take a minute to complete.

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 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 main menu () , select **Cloud Resources**.
 - 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 Click **Keep** to preserve your changes.

The operation can take a minute to complete.

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 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 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 vCloud 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 vCloud 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 edge gateway in a vCloud 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 vCloud 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 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 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 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 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 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 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 Edge Gateway](#)

By default, the system sets some SSL VPN-Plus settings on an edge gateway in your vCloud Director environment. You can use the **General Settings** screen on the **SSL VPN-Plus** tab in the vCloud 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 edge gateway.

Procedure

- 1 Open Edge Gateway Services.
 - a From the main menu () , select **Cloud Resources**.
 - 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 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 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 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 Edge Gateway](#).

Create an IP Pool for Use with SSL VPN-Plus on an 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.

Option	Action
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.
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 Edge Gateway](#).

Add a Private Network for Use with SSL VPN-Plus on an 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 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 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 Edge Gateway Firewall Rule](#).

Configure an Authentication Service for SSL VPN-Plus on an 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 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 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 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 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 login	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 Edge Gateway

By default, the system sets some SSL VPN-Plus settings on an edge gateway in your vCloud Director environment. You can use the **General Settings** screen on the **SSL VPN-Plus** tab in the vCloud 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 edge gateways in a vCloud 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 (AES265-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. You can view the size of your edge gateway by logging in to the vCloud Director Web console, navigating to the edge gateway, and using the **Properties** action to view the edge gateway configuration. See the *vCloud Director Administrator's Guide* for information about using the vCloud Director Web console.

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 edge gateway.

2 [Configure the IPsec VPN Site Connections for the Edge Gateway](#)

Use the **IPsec VPN Sites** screen in the vCloud 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 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 edge gateway.

Procedure

- 1 Open Edge Gateway Services.
 - a From the main menu () , select **Cloud Resources**.
 - 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 Edge Gateway](#).

Configure the IPsec VPN Site Connections for the Edge Gateway

Use the **IPsec VPN Sites** screen in the vCloud 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 vCloud 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, 192.168.99.0/24.</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> <ul style="list-style-type: none"> ■ <code>passthroughSubnets=PeerSubnet/IPAddress</code> to support overlapping subnets.

4 Click **Keep**.

5 Click **Save changes**.

The save operation can take a minute to complete.

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 Edge Gateway](#).

Enable the IPsec VPN Service on an 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 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

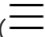
Use the **Global Configuration** screen to configure IPsec VPN authentication settings at an edge gateway level. On this screen, you can set a global pre-shared key and enable certification authentication.

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 main menu () , select **Cloud Resources**.
 - 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 edge gateways in a vCloud Director environment support L2 VPN. L2 VPN allows extension of your organization virtual data center by allowing 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.

The NSX software provides the L2 VPN capabilities of an edge gateway. L2 VPN allows you to 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. See the *vCloud Director Administrator's Guide* for the steps on creating an external routed organization virtual data center network.

Procedure

1 Navigate to the L2 VPN Screen

To begin configuring the L2 VPN service for an edge gateway, you must navigate to the **L2 VPN** screen.

2 Configure the 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 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 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 edge gateway, you must navigate to the **L2 VPN** screen.

Procedure

1 Open Edge Gateway Services.

- a From the main menu () , select **Cloud Resources**.
- 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 Edge Gateway as an L2 VPN Server](#).

Configure the 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. See the *vCloud Director Administrator's Guide* for the steps on creating an external routed organization virtual data center network.
- [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	<p>Select at least one subinterface to be stretched with the client.</p> <p>The subinterfaces available for selection are those organization virtual data center networks configured as subinterfaces on the edge gateway.</p>
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**.

The save operation can take a minute to complete.

What to do next

Enable the L2 VPN service on this edge gateway. See [Enable the L2 VPN Service on an Edge Gateway](#).

Configure the 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.

Option	Description
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.
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**.

The save operation can take a minute to complete.

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 Password	Enter the proxy server authentication credentials.

6 To enable server certification validation, click **Change CA certificate** and select the appropriate CA certificate.

7 Click **Save changes**.

The save operation can take a minute to complete.

What to do next

Enable the L2 VPN service on this edge gateway. See [Enable the L2 VPN Service on an Edge Gateway](#).

Enable the L2 VPN Service on an 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 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 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 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 vCloud 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 vCloud 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 vCloud 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 main menu () , select **Cloud Resources**.
 - 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 vCloud 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	Type the key type, either RSA or DSA, for the certificate. RSA is typically used. The key type defines the encryption algorithm for communication between the hosts. 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](#).

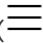
Import the CA-Signed Certificate Corresponding to the CSR Generated for an Edge Gateway

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 main menu () , select **Cloud Resources**.
 - 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 main menu () , select **Cloud Resources**.
 - 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.

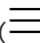
You usually add the root certificate of your company or organization as a CA certificate. A typical use is for SSL VPN, where you want to authenticate VPN clients using certificates. Client certificates can be distributed to the VPN clients and when the VPN clients connect, their client certificates 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 main menu () , select **Cloud Resources**.
 - 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 main menu () , select **Cloud Resources**.
 - 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.

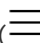
Add a Service Certificate to the Edge Gateway

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 main menu () , select **Cloud Resources**.
 - 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.
 - 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) Type 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 vCloud 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 main menu () , 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 main menu () , 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.

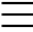
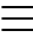
Create a MAC Set for Use in Firewall Rules

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"> a From the main menu () , 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 main menu () , 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 **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 main menu () , 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 main menu () , 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 main menu () , 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 main menu () , 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.

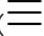
Results

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 main menu () , select **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

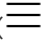
Activate or Deactivate Distributed Routing on an Edge Gateway

After you activate vCloud 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 main menu () , select **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**.

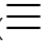
Modify the External Networks and the Edge Gateway Settings

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 Edge Gateway](#).

To modify the distributed routing setting, see [Activate or Deactivate Distributed Routing on an Edge Gateway](#).

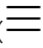
Procedure

- 1 From the main menu () , select **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, enable or disable FIPS mode and high availability state, and change the edge gateway size configuration.

Procedure

- 1 From the main menu () , select **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.


Option	Description
Compact	Requires less memory and fewer compute resources.
Large	Provides increased capacity and performance than with the Compact option. 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 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.

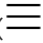
Procedure

- 1 From the main menu () , select **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 main menu () , select **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.

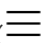
If you do not enter an IP address for a network, the system assigns an arbitrary IP address to this network.
- 5 To confirm the changes, click **Save**.

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. vCloud 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 vCloud Director.

Procedure

- 1 From the main menu () , select **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

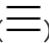
The system combines overlapping IP ranges.

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 main menu () , select **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.

You can redeploy legacy edge gateways to migrate the edge gateways to newly created edge clusters.

When you redeploy an edge gateway, vCloud Director deletes it and recreates it with the latest configurations.

Procedure

- 1 From the main menu () , select **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 main menu () , select **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 main menu () , select **Cloud Resources**.
 - 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

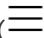
You can enable logging for an edge gateway. In addition to enabling logging for the features for which you want to collect log data, to complete the configuration, you must have a Syslog server to receive the collected log data. When you configure a Syslog server on the Edge Settings screen, you are 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 main menu () , select **Cloud Resources**.
- 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 vCloud Director system administrator has already configured a Syslog server for the vCloud 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 main menu () , select **Cloud Resources**.
 - 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 SSH access to this edge gateway.
Password	By default, the SSH username 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 SSH access to this edge gateway.

Managing Organization Virtual Data Center Networks



This chapter includes the following topics:

- [Managing NSX-T Organization Virtual Data Center Networks](#)

Managing NSX-T Organization Virtual Data Center Networks

Only system administrators can create, modify, and delete organization virtual data center networks that are based on NSX-T logical switches.

To manage organization virtual data center networks, system administrators must log in to the Service Provider Admin Portal and navigate to the vCloud Director Tenant Portal of the target organization.

For information about managing organization virtual data center networks that are based on NSX Data Center for vSphere, see *vCloud Director Tenant Portal Guide*.

Add an NSX-T Organization Virtual Data Center Network

As a system administrator, you can create an organization virtual data center network by importing a logical switch from an associated NSX-T Manager instance.

Note With an NSX-T logical switch, you can create only an IPv4 isolated organization network. You cannot create a direct or routed organization network based on an NSX-T logical switch.


Prerequisites

- The provider virtual data center that backs the target organization virtual data center must be associated with an NSX-T Manager instance.
- You created at least one NSX-T logical switch that is not in use by other organization virtual data center networks.

For information about configuring NSX-T logical switches, see the *NSX-T Administration Guide*. For information, about creating a Provider VDC backed by an NSX-T Manager instance, see *vCloud API Programming Guide for Service Providers*.

Procedure

- 1 Navigate to the vCloud Director Tenant Portal of the target organization.

- a From the main menu () , select **Cloud Resources**.

- b Under **Organizations**, click the name of the target organization.

You are redirected to **Datacenters** view of the vCloud Director Tenant Portal for this organization.

- 2 If there are multiple VDCs in the organization, click the card of the target organization VDC.

- 3 In the left panel, under **Networks**, click **Network**.

- 4 Click **Import**.

The **Import Logical Switch** wizard appears.

- 5 Enter a name and, optionally, a description for the new organization VDC network, and click **Next**.

- 6 From the list of available NSX-T logical switches, select the target switch by clicking the radio button next to the switch name, and click **Next**.

- 7 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**.

If the switch is configured with a subnet, this information is prepopulated.

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

You can add multiple IP addresses and IP ranges.

- 9 Click **Next**.


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

Edit an NSX-T Organization Virtual Data Center Network

You can modify the name, description, DNS settings, and static IP pool of an organization virtual data center network that is based on an NSX-T logical switch. You cannot edit the network Classless Inter-Domain Routing (CIDR) settings.

Procedure

- 1 Navigate to the vCloud Director Tenant Portal of the target organization.

- a From the main menu () , select **Cloud Resources**.

- b Under **Organizations**, click the name of the target organization.

You are redirected to **Datacenters** view of the vCloud Director Tenant Portal for this organization.

- 2 If there are multiple VDCs in the organization, click the card of the target organization VDC.
- 3 In the left panel, under **Networks**, click **Network**.
- 4 Click the radio button next to the name of the target network and click **Modify**.

The **Edit Org VDC Network** wizard appears.

- 5 (Optional) On the **General** tab, edit the name and the description of the network.
- 6 (Optional) On the **Configure Network** tab, edit the DNS settings, and static IP pool of the network.

You can add, modify, and remove IP addresses and IP ranges.


- 7 Click **Save**.

Delete an NSX-T Organization Virtual Data Center Network

If you no longer use an NSX-T organization virtual data center network, you can delete this network.

Procedure

- 1 Navigate to the vCloud Director Tenant Portal of the target organization.

- a From the main menu () , select **Cloud Resources**.
- b Under **Organizations**, click the name of the target organization.

You are redirected to **Datacenters** view of the vCloud Director Tenant Portal for this organization.

- 2 If there are multiple VDCs in the organization, click the card of the target organization VDC.
- 3 In the left panel, under **Networks**, click **Network**.
- 4 Click the radio button next to the name of the target network and click **Delete**.
- 5 To confirm, click **OK**.

Managing SDDCs and SDDC Proxies

9

Starting with version 9.7, vCloud Director can act as an HTTP proxy server between tenants and the underlying vSphere environment. A Software-Defined Data Center (SDDC) encapsulates the infrastructure of an attached vCenter Server instance. An SDDC proxy is an access point to a component from an SDDC, for example, a vCenter Server instance, an ESXi host, or an NSX Manager instance.

With the SDDC feature, you can use vCloud 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 SDDC only to its organization. The tenant does not share these resources with other tenants. The tenant can access this SDDC by using a UI or API proxy without a VPN required.
- You can use vCloud Director as a lightweight directory to register all your vCenter Server instances.
- You can use vCloud Director as an API endpoint for all your vCenter Server instances.

Before you create an SDDC, you must attach the target vCenter Server instance to vCloud Director. See [Attach a vCenter Server Instance Alone or Together with an NSX Manager Instance](#).

Note By default, with an attached vCenter Server instance, you can create either a provider VDC or an SDDC. If you created a provider VDC backed by an vCenter Server instance, you cannot use this vCenter Server instance to create an SDDC, and the reverse. You can use the vCloud API to modify the system settings of your vCloud Director installation so that a vCenter Server instance can back both a provider VDC and an SDDC.

You can create and publish SDDCs and SDDC proxies to organizations in your cloud. Users can use the SDDC proxies to access the underlying vSphere environment. Users can log in to the UI or API of the proxied components by using their vCloud Director accounts.

SDDCs in vCloud Director remove the requirement for vCenter Server to be publically accessible. To control the access, you can enable and disable an SDDC in vCloud Director, and you can enable and disable an SDDC proxy.

Creating and Managing SDDCs and SDDC Proxies

To create and manage SDDCs and proxies, you must use the vCloud OpenAPI. See *Getting Started with vCloud OpenAPI* at <https://code.vmware.com>.

Important vCloud Director requires a direct network connection to each vCenter Server instance for use as an SDDC. If the vCenter Server instance uses an external Platform Services Controller instance, vCloud Director requires a direct network connection to the Platform Services Controller instance as well.

To use VMware OVF Tool in a proxied SDDC, vCloud Director requires a direct connection to each ESXi host.

- 1 Create an SDDC backed by an attached and enabled vCenter Server instance.
vCloud Director creates the SDDC with a default proxy for the vCenter Server instance. If the vCenter Server instance uses an external Platform Services Controller instance, vCloud Director creates a proxy for the Platform Services Controller instance as well.
- 2 Get the certificate and the thumbprint of the created proxies, and verify that the certificate and the thumbprint are present and correct.
- 3 Enable the SDDC.
- 4 Publish the SDDC to one or more organizations.
- 5 To enable users to access the SDDCs and the SDDC proxies from the vCloud Director Tenant Portal, you must publish the **CPOM extension** plug-in to their organizations. See [Publish or Unpublish a Plug-in from an Organization](#).

After you create and publish an SDDC, you can add, edit, enable, disable, and remove its SDDC proxies.

Note When you add a proxy to an SDDC, 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.

Managing System Administrators and Roles

10

By using the vCloud Director Web Console, you can add system administrators to vCloud 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 vCloud Director 9.5, service providers can create provider roles and manage provider users and groups by using the vCloud Director Service Provider Admin Portal or by using the vCloud OpenAPI. For information about managing provider roles, users, and groups, see the *vCloud 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 vCloud Director. A role associates a role name with a set of rights. Each organization can have different rights and roles.

vCloud 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 vCloud 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.

vCloud Director 9.5 introduces rights bundles and global tenant roles which system administrators can use to manage the rights and roles that are available to each organization.

After you install vCloud 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 vCloud 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 vCloud 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 vCloud 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 vCloud 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 *vCloud Director Tenant Portal Guide*.

Predefined Roles and Their Rights

Each vCloud 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. 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 vCloud API 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 use the vCloud Director Web Console, tenant portal, or vCloud OpenAPI to 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 the user is defined by an OAuth Identity Provider, the user is assigned the roles named in the `roles` array of the user's OAuth token.
- If the user is defined by a SAML Identity Provider, the user is assigned the roles named in the SAML attribute whose name appears in the `RoleAttributeName` element, which is in the `SamlAttributeMapping` element in the organization's `OrgFederationSettings`.

If a user is assigned the **Defer to Identity Provider** role but no matching role or group name is available in your organization, the user can log in to the organization but has no rights. If an Identity Provider associates a user with a system-level role such as **System Administrator**, the user can log in to the organization but has no rights. You must manually assign a role to such users.

Except the **Defer to Identity Provider** role, each predefined role includes a set of default rights. Only a **System Administrator** can modify the rights in a predefined role. If a **System administrator** modifies a predefined role, the modifications propagate to all instances of the role in the system.

Rights in Predefined Global Tenant Roles

Various rights are common to multiple predefined global roles. These rights are granted by default to all new organizations, and are available for use in other roles created by the **Organization Administrator**.

Table 10-1. Rights Included in the Global Tenant Roles in vCloud Director

Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
Catalog: Add a vApp from My Cloud	X	X	X		
Catalog: Allow External Publishing / Subscriptions for the Catalogs	X	X			
Catalog: Change Owner	X				
Catalog: Create / Delete a Catalog	X	X			
Catalog: Edit Catalog Properties	X	X			
Catalog: Share a Catalog to Other Organizations	X	X			
Catalog: Share a Catalog to Users / Groups within Current Organization	X	X			
Catalog: View Private and Shared Catalogs within Current Organization	X	X	X		
Catalog: View Shared Catalogs from Other Organizations	X				
Catalog Item: Add to My Cloud	X	X	X	X	
Catalog Item: Copy / Move a vApp Template / Media	X	X	X		
Catalog Item: Create / Upload a vApp Template / Media	X	X			
Catalog Item: Edit vApp Template / Media	X	X			
Catalog Item: Enable vApp Template / Media Download	X	X			

Table 10-1. Rights Included in the Global Tenant Roles in vCloud Director (continued)

Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
Catalog Item: View vApp Templates / Media	X	X	X	X	
Custom Entity: View All Custom Entity Instances in Organization	X				
Custom Entity: View Custom Entity Instance	X				
Disk: Change Owner	X	X			
Disk: Create a Disk	X	X	X		
Disk: Delete a Disk	X	X	X		
Disk: Edit Disk Properties	X	X	X		
Disk: View Disk Properties	X	X	X	X	
Distributed Firewall: Configure Distributed Firewall Rules	X				
Distributed Firewall: Enable / Disable Distributed Firewall	X				
Distributed Firewall: View Distributed Firewall Rules	X				
Edge Cluster: View Edge Cluster	X				
Edge Cluster: Manage Edge Cluster	X				
Gateway: Configure Syslog Server	X				
Gateway: Configure System Logging	X				
Gateway: Convert to Advanced Gateway	X				
Gateway: View Gateway	X				
Gateway: Enable Distributed Routing	X				
Gateway: Import Edge Gateway	X				
Gateway Services: BGP Routing Configure					
Gateway Services: DHCP Configure	X				
Gateway Services: Firewall Configure	X				
Gateway Services: IPSEC VPN Configure	X				
Gateway Services: L2 VPN Configure					
Gateway Services: Load Balancer Configure	X				

Table 10-1. Rights Included in the Global Tenant Roles in vCloud Director (continued)

Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
Gateway Services: NAT Configure	X				
Gateway Services: OSPF Routing Configure	X				
Gateway Services: Remote Access Configure	X				
Gateway Services: SSL VPN Configure	X				
Gateway Services: Static Routing Configure	X				
Gateway Services: BGP Routing View Only	X				
Gateway Services: DHCP View Only	X				
Gateway Services: Firewall View Only	X				
Gateway Services: IPSEC VPN View Only	X				
Gateway Services: L2 VPN View Only	X				
Gateway Services: Load Balancer View Only	X				
Gateway Services: NAT View Only	X				
Gateway Services: OSPF Routing View Only	X				
Gateway Services: Remote Access View Only	X				
Gateway Services: SSL VPN View Only	X				
Gateway Services: Static Routing View Only	X				
General: Administrator Control	X				
General: Administrator View	X				
General: Send Notification	X				
Hybrid Tunnel: Acquire Control Ticket	X				
Hybrid Tunnel: Acquire From-the-Cloud Tunnel Ticket	X				
Hybrid Tunnel: Acquire To-the-Cloud Tunnel Ticket	X				
Hybrid Tunnel: Create From-the-Cloud Tunnel	X				

Table 10-1. Rights Included in the Global Tenant Roles in vCloud Director (continued)

Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
Hybrid Tunnel: Create To-the-Cloud Tunnel	X				
Hybrid Tunnel: Delete From-the-Cloud Tunnel	X				
Hybrid Tunnel: Delete To-the-Cloud Tunnel	X				
Hybrid Tunnel: Update From-the-Cloud Tunnel Endpoint Tag	X				
Hybrid Tunnel: View the Cloud Tunnel Server Settings	X				
Hybrid Tunnel: View From-the-Cloud Tunnel	X				
Hybrid Tunnel: View To-the-Cloud Tunnel	X				
Organization: Allow Access to All Organization VDCs	X				
Organization: Edit Access Control List of Organization VDCs	X				
Organization: Edit Federation Settings	X				
Organization: Edit Leases Policy	X				
Organization: Edit Organization Associations	X				
Organization: Edit Organization Network Properties	X				
Organization: Edit Organization OAuth Settings	X				
Organization: Edit Organization Properties	X				
Organization: Edit Password Policy	X				
Organization: Edit Quotas Policy	X				
Organization: Edit SMTP Settings	X				
Organization: Implicitly Import User/ Group from IdP while Editing VDC ACL	X				
Organization: View Access Control List of Organization VDCs	X				
Organization: View Catalog ACL	X	X			
Organization: View Organization Networks	X				

Table 10-1. Rights Included in the Global Tenant Roles in vCloud Director (continued)

Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
Organization: View Organizations	X	X	X		
Organization: View vApp ACL	X	X	X	X	
Organization VDC: Edit Organization VDC Name and Description	X				
Organization VDC: Edit VM-VM Affinity Rule	X	X	X		
Organization VDC: Edit Organization VDC Extended Properties	X				
Organization VDC: Manage Firewall	X				
Organization VDC: Set Default Storage Policy	X				
Organization VDC: View Compute Policies for an Organization VDC	X	X	X	X	
Organization VDC: View Organization VDC Extended Properties	X				
Organization VDC Network: View Properties	X				
Organization VDC Network: Edit Properties	X				
Organization VDC Network: Import Network	X				
Organization VDC: View Organization VDCs	X				
Organization VDC Template: Instantiate Organization VDC templates	X				
Organization VDC Template: View VDC templates	X				
Provider Network: View Provider Network	X				
Provider Network: Create / Delete Provider Network	X				
Role: Create / Update / Delete a Role	X				
Service Library: View Services Making Up the Service Library	X				
User: View Group / User	X				
VCD Extension: View Tenant Portal Plugin Information	X	X	X	X	
VDC Group: View VDC Group	X				

Table 10-1. Rights Included in the Global Tenant Roles in vCloud Director (continued)

Right Name	Organization Administrator	Catalog Author	vApp Author	vApp User	Console Access Only
VDC Group: Configure VDC Group	X				
VM Monitoring: View historic metrics for the Organization	X				
VM Monitoring: View historic metrics for the Organization VDC	X				
vApp: Access to VM Console	X	X	X	X	X
vApp: Allow Metadata Mapping Domain to vCenter Server	X	X	X		
vApp: Change Owner	X				
vApp: Change vApp Template Owner	X	X			
vApp: Copy a vApp	X	X	X	X	
vApp: Create / Reconfigure vApp	X	X	X		
vApp: Create / Revert / Remove / a Snapshot	X	X	X	X	
vApp: Delete a vApp	X	X	X	X	
vApp: Download a vApp	X	X	X		
vApp: Edit / View VM Boot Options	X	X	X		
vApp: Edit VM CPU	X	X	X		
vApp: Edit VM Hard Disk	X	X	X		
vApp: Edit VM Memory	X	X	X		
vApp: Edit VM Network	X	X	X	X	
vApp: Edit VM Properties	X	X	X	X	
vApp: Edit vApp Properties	X	X	X	X	
vApp: Edit VM Compute Policy	X	X	X		
vApp: Manage VM Password Settings	X	X	X	X	X
vApp: Share a vApp	X	X	X	X	
vApp: Start / Stop / Suspend / Reset a vApp	X	X	X	X	
vApp: Upload a vApp	X	X	X		
vApp: View VM metrics	X		X	X	

For information about the new rights that vCloud Director 9.7 introduces, see [New Rights in This Release](#).

New Rights in This Release

vCloud Director 9.7 introduces new rights, which you might want to add to any existing global roles that you published to your tenants.

Right	Description	Default Role
SDDC: View SDDC	Allows you to view all SDDCs that are published to your organization. The system administrator can view all SDDCs.	System Administrator and Organization Administrator
SDDC: Manage SDDC	Allows you to add, remove, and edit SDDCs.	System Administrator
SDDC: Manage SDDC Proxy	Allows you to add, remove, enable, and disable SDDC proxies.	System Administrator
Service Applications: View Service Applications	Allows you to see the list of registered service applications. Used for VMC accounts.	System Administrator
Service Applications: Register VMC SDDC	Allows you to create, view, edit, and remove service applications. Used for VMC accounts.	System Administrator
Service Applications: Manage Service Applications	Allows you to register service applications. Used for VMC accounts.	System Administrator
Edge Cluster: View Edge Cluster	Allows you to see a list of edge cluster, and to retrieve an individual edge cluster.	System Administrator and Organization Administrator
Edge Cluster: Manage Edge Cluster	Allows you to create, edit, and remove edge clusters.	System Administrator and Organization Administrator
vApp: Edit VM Compute Policy	Allows users to change the compute policy of a virtual machine.	system administrator , organization administrator , Catalog Author , and vApp Author
Gateway: Import Edge Gateway	Allows you to import a Tier-1 router as an edge gateway.	System Administrator and Organization Administrator

For information about managing rights and roles, see the *vCloud Director Service Provider Admin Portal Guide*.


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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 vCloud 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](#).

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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 vCloud 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 vCloud 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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 vCloud 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 vCloud 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](#).


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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 main menu () , select **Administration**.
- 2 In the left panel, under **Tenant Access Control**, click **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 vCloud Director Provider organization.

For information about managing tenant roles, see the *vCloud Director Tenant Portal Guide*.


Create a Provider Role

You can create a role in your vCloud Director Provider organization.

After the initial vCloud 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 main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 vCloud 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**.


Results

The newly created role is available for assigning to users in your Provider organization.

View or Edit a Provider Role

You can view the rights that are included in a role that is local to your vCloud Director Provider organization. You can modify the name, the description, and the rights of a role.

Procedure

- 1 From the main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 vCloud Director Provider organization.

Prerequisites

The role that you want to delete must not be assigned to any user.

Procedure

- 1 From the main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 vCloud Director Provider organization.

For information about managing organization users and groups, see the *vCloud 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 *vCloud Director Tenant Portal Guide*.

Create a Provider User

You can create a user in your vCloud Director Provider organization.

During the vCloud 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 main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 vCloud 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 main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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.

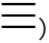
Procedure

- From the main menu () , select **Administration**.
- In the left panel, under **Provider Access Control**, click **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 vCloud Director.

Procedure

- From the main menu () , select **Administration**.
- In the left panel, under **Provider Access Control**, click **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

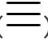
You can remove a user from your vCloud Director Provider organization by deleting the user account.

To delete a stranded user that lost access to the system because their LDAP group was deleted, use the vCloud Director API.

Prerequisites

Deactivate the user that you want to delete. See [Activate or Deactivate a Provider User](#).

Procedure


- 1 From the main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 the *vCloud Director Administrator's Guide*.

Procedure

- 1 From the main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 *vCloud Director Tenant Portal Guide*.


Import a Provider Group

You can import groups to your vCloud 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 main menu () , select **Administration**.

- 2 In the left panel, under **Provider Access Control**, click **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 vCloud Director Provider organization.

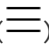
Procedure

- 1 From the main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 vCloud Director Provider organization

Procedure

- 1 From the main menu () , select **Administration**.
- 2 In the left panel, under **Provider Access Control**, click **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 vCloud Director system administrator can control system-wide settings related to LDAP, email notification, licensing, and general system preferences.

This chapter includes the following topics:

- [Managing Identity Providers](#)
- [Managing Plug-Ins](#)
- [Customizing the vCloud Director Portals](#)

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 vCloud 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.

Configure a System LDAP Connection

To provide vCloud 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 main menu () , 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **Organizations**.
- 3 Click the name of the target organization.
You are redirected to the vCloud Director Tenant Portal of the organization.
- 4 From the main menu () , 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 vCloud 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 a system or organization 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 vCloud Director with the LDAP server at any time.

Prerequisites

If you plan to connect to an LDAPS server, verify that you have a properly constructed certificate for the improved LDAP support in Java 8 Update 181. 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.
Base distinguished name	The base distinguished name (DN) is the location in the LDAP directory where vCloud Director to connect. To connect at the root, enter only the domain components, for example, DC=example,DC=com . To connect to a node in the tree, enter the distinguished name for that node, for example, OU=ServiceDirector,DC=example,DC=com . Connecting to a node limits the scope of the directory available to vCloud Director.
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 icon () and import an LDAP SSL certificate or select Accept all certificates .
Authentication method	Simple authentication consists of sending the LDAP server the user's DN and password. If you are using LDAP, the LDAP password is sent over the network in a clear text. If you want to use Kerberos, you must configure the LDAP connection by using the vCloud Director Web Client. For more information, see the <i>vCloud Director Administrator's Guide</i> .
User name	The full LDAP DN user name for connecting to the LDAP server. If anonymous read support is enabled on your LDAP server, you can leave these text boxes blank.
Password	The password for connecting to the LDAP server. If anonymous read support is enabled on your LDAP server, you can leave these text boxes blank.

- 2 Click the **User Attributes** tab, examine the default values for the user attributes, and, if your LDAP directory uses different schema, modify the values.
- 3 Click the **Group Attributes** tab, examine the default values for the group attributes, and, if your LDAP directory uses different schema, modify the values.
- 4 Click **Save**.

5 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**.

6 To synchronize vCloud Director with the configured LDAP server, click **Sync**.

vCloud 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 vCloud 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.

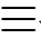
If an imported user or group is assigned the **Defer to Identity Provider** role, the roles are assigned based on the information gathered from the **Roles** attribute in the token. If a different attribute is used, this attribute name can be configured using API and only the **Roles** attribute is configurable. 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.


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 main menu () , 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 vCloud 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**.

Results

Managing Plug-Ins

vCloud Director plug-ins extend the functions of the Service Provider Admin Portal and the vCloud 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 vCloud Director.

CPOM extension

Provides the capability for viewing and managing SDDCs and SDDC proxies by using the vCloud Director Tenant Portal.

Customize Portal

Provides the capability for customizing the vCloud Director Service Provider Admin Portal and the vCloud Director Tenant Portal.

vCloud Availability

The VMware vCloud[®] Availability[™] plug-in provides the capability to access vCloud Availability Portal directly from the vCloud Director user interface. For more information, see [vCloud Availability Documentation](#).


Upload a Plug-in

You can upload additional plug-ins to your vCloud 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 main menu () , select **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 vCloud Director Service Provider Admin Portal.
Tenants	The plug-in function becomes available in the vCloud 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 main menu () , select **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 vCloud Director Service Provider Admin Portal.

Procedure

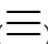
- 1 From the main menu () , select **Customize Portal**.
- 2 Select the check boxes next to the names of the plug-ins that you want to remove, and click **Delete**.
- 3 To confirm, click **Save**.

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 main menu () , select **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 vCloud Director Service Provider Admin Portal.
Tenants	The plug-in function becomes available in the vCloud 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 vCloud 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 vCloud Director Service Provider Admin Portal and for the vCloud Director Tenant Portal of each organization. In addition, you can modify and add custom links to the two upper right menus in the vCloud Director portals.

Note To customize your branding attributes and links, you must use the `branding` vCloud OpenAPI methods. See *Getting Started with vCloud OpenAPI* at <https://code.vmware.com>.

Portal Branding

As part of the installation, vCloud 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 vCloud Director Service Provider Admin Portal. The vCloud 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 *vCloud Director Administrator's 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.

vCloud Director supports the following custom variables in the `url` value for a custom link:

Table 11-1. 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}/vdcs
```

in the vCloud Director Tenant Portal for organization myorg is converted to:

```
url: https://host:port/tenant/myorg/vdcs
```

Monitoring vCloud Director

12

System administrators can monitor completed and in-progress operations and view resource usage information at the provider virtual datacenter, organization virtual datacenter, and datastore level.

This chapter includes the following topics:

- [vCloud Director and Cost Reporting](#)
- [View Use Information for a Provider Virtual Data Center](#)

vCloud Director and Cost Reporting

You can use VMware vRealize Operations Tenant App for vCloud Director to configure a cost reporting system for vCloud 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 vCloud 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 vCloud Director as a Service Provider* and *Using vRealize Operations Tenant App for vCloud 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 main menu () , select **Cloud Resources**.
- 2 In the left panel, click **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.

Managing Services

13

The Content Libraries view in the vCloud 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 vCloud 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 vCloud Director

You integrate vRealize Orchestrator with vCloud Director through the vCloud Director Service Provider Admin Portal.

Integrating vRealize Orchestrator with vCloud Director extends the base functionality of vCloud Director by allowing service provider administrators to develop complex automation tasks through workflow orchestration and utilization of third-party plug-ins.

Through the vCloud Director Service Provider Admin Portal, service provider administrators are able to view, import, and execute workflows from registered vRealize Orchestrator server instances.

In the vCloud 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 vCloud Director


To leverage orchestration of workflows and automation of tasks through vRealize Orchestrator in vCloud Director, you register a vRealize Orchestrator instance in the vCloud 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 vCloud 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 main menu (), select **Content 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 the  button.
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 8281. Note vCloud Director connects to the API interface of vRealize Orchestrator.
Username	A user account that is member of the vRealize Orchestrator administrators group.
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.

- 4 Click **OK** to complete the registration.
The vRealize Orchestrator server is registered with vCloud Director.

Create a Service Category

You can organize services in service categories.

Procedure



- 1 From the main menu (), select **Content Libraries**
 - a From the left panel, select **Service Management**.
 - b Navigate to the **Service Categories** tab.A list of existing server categories appears.
- 2 To create a new service category, click the  button.
The **New Service Category** dialog appears.
- 3 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

- 1 From the main menu (), select **Content Libraries**
 - a From the left panel, select **Service Management**.
 - b Navigate to the **Service Categories** tab.A list of existing server categories appears.
- 2 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 vCloud Director.

Prerequisites

- Register a vRealize Orchestrator instance. See [Register a vRealize Orchestrator Instance with vCloud Director](#).
- Create a service category. See [Create a Service Category](#).

Procedure

- 1 From the main menu () , select **Content Libraries**.
 - a From the left panel, select **Service Library**.

Available services display in a card view of twelve 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 main menu () , select **Content Libraries**.
 - a From the left panel, select **Service Library**.

Available services display in a card view of twelve 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 main menu () , select **Content Libraries**.

- a From the left panel, select **Service Library**.

Available services display in a card view of twelve 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 vCloud Director service, vCloud 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 vCloud Director
_vcd_apiEndpoint	vCloud Director REST API endpoint

Change a Service Category

You can change the category, to which a service belongs.

Procedure

- 1 From the main menu () , select **Content Libraries**.
 - a From the left panel, select **Service Library**.

Available services display in a card view of twelve 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 main menu () , select **Content Libraries**.
 - a From the left panel, select **Service Library**.

Available services display in a card view of twelve 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 main menu () , select **Content Libraries**.
 - a From the left panel, select **Service Library**.

Available services display in a card view of twelve 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 Custom Entities

14

The custom entity definitions in vCloud 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 vCloud 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.


This chapter includes the following topics:

- [Search for a Custom Entity](#)
- [Edit a Custom Entity Definition](#)
- [Add a Custom Entity Definition](#)
- [Custom Entity Instances](#)
- [Associate an Action to a Custom Entity](#)
- [Dissociate an Action From a Custom Entity](#)
- [Publish a Custom Entity](#)
- [Delete a Custom Entity](#)

Search for a Custom Entity

You can search for a custom entity by its name.

Procedure

- 1 From the main menu () , select **Content Libraries**.
 - a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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 main menu () , select **Content Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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 main menu () , select **Content Libraries**.


- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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 Click the  icon to add a new custom entity.

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 vCloud Director shows the output parameter as an instance of a custom entity.


Procedure

- From the main menu () , select **Content Libraries**.
 - From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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.

- In the card of the selected custom entity, click **Intances**.

The available instances display in a grid view.


- Click the list bar () on the left of each entity to display the associated workflows.

Clicking on a workflow initiates a workflow run which takes the entity instance as an input parameter.

Associate an Action to a Custom Entity

By associating an action to a custom entity definition, you can execute a set of vRealize Orchestrator workflows on the instances of a particular custom entity.

Procedure

- 1 From the main menu () , select **Content Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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 main menu () , select **Content Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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 main menu () , select **Content Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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 main menu () , select **Content Libraries**.

- a From the left panel, select **Custom Entity Definitions**.

The list of custom entities displays in a card view of twelve 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.