

VMware Integrated OpenStack Administrator Guide

Modified on 14 Nov 2017

VMware Integrated OpenStack 4.0



vmware®

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

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

If you have comments about this documentation, submit your feedback to

docfeedback@vmware.com

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2015–2017 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About This Book	6
Updated Information	7
1 About VMware Integrated OpenStack	8
Internationalization	8
OpenStack Foundation Compliance	9
VMware Integrated OpenStack System Requirements	9
OpenStack Instances in vSphere Web Client	13
Monitor OpenStack Instances in the vSphere Web Client	16
Customer Experience Improvement Program	16
2 Managing Your VMware Integrated OpenStack Deployment	18
Managing Your Deployment Configuration	18
Managing Your Network Configuration	28
Adding Capacity in the vSphere Web Client	37
Configure the Backup Service for Block Storage	40
Back Up the VMware Integrated OpenStack Deployment	42
Restore VMware Integrated OpenStack from a Backup	43
Failure Recovery	45
VMware Integrated OpenStack Log File Locations	47
Upgrade to VMware Integrated OpenStack 4.0	49
Updating Your VMware Integrated OpenStack Deployment	55
Customize the Dashboard Logos and Background	59
Use Profiling to Trace OpenStack Deployments	63
Configure NUMA for Use With VMware Integrated OpenStack	65
3 Managing OpenStack Projects and Users	67
Create an OpenStack Project	67
Modify a Project	68
Working with Security Groups	69
Create a Cloud User Account in OpenStack	76
Modify a User Account	77
4 Working with Instances in OpenStack	78
Import vSphere VMs into VMware Integrated OpenStack	78
Create a Snapshot from an Instance	83
Control the State of an Instance	83

- Track Instance Use 83
- Use DRS to Control OpenStack Instance Placement 84
- Using Affinity and Anti-Affinity to Place OpenStack Instances 88
- Apply QoS Resource Allocation to Existing Instances 90
- Define Default Nova Storage for OpenStack Instances 91

- 5 Working with Volumes and Volume Types in OpenStack 93**
 - Modify the Default Cinder Volume Adapter Type 93
 - Create a Volume Type 94
 - Delete a Volume Type 95
 - Migrating Volumes Between Datastores 96

- 6 Managing Images for the Image Service 100**
 - Import Images to the Image Service 100
 - Modify Image Settings 106
 - Modify Image Resource Metadata 106
 - Configuring Images for Windows Guest Customization 107
 - Configure QoS Resource Allocation for Instances Using Image Metadata 109
 - Delete an Existing Image 112
 - Migrating Images 112
 - Add a VM Template as an Image 114
 - Configure Images to Enable Live Resize of VMs Deployed From That Image 116
 - Modify the Default Behavior for Nova Snapshots 117
 - Modify the Default Cinder Upload-to-Image Behavior 117

- 7 Working with Flavors 119**
 - Default Flavor Configurations 119
 - Create a Flavor 119
 - Delete a Flavor 120
 - Modify Flavor Metadata 121
 - Configure QoS Resource Allocation for Instances Using Flavor Metadata 122

- 8 Working with VMware Integrated OpenStack Carrier Edition 125**
 - Use Tenant Virtual Data Center for Secure Multi-Tenancy and Resource Allocation 125
 - Configure Passthrough Devices for Instances 128

- 9 VMware Integrated OpenStack CLI Command Reference 134**
 - viocli backup Command 134
 - viocli dbverify Command 135
 - viocli deployment Command 135
 - viocli ds-migrate-prep Command 137
 - viocli epops Command 137

- [viocli inventory-admin Command](#) 138
- [viocli recover Command](#) 140
- [viocli restore Command](#) 141
- [viocli rollback Command](#) 141
- [viocli services Command](#) 142
- [viocli show Command](#) 142
- [viocli upgrade Command](#) 143
- [viocli volume-migrate Command](#) 144

10 vRealize Automation Integration 146

- [Integrate VMware Integrated OpenStack with vRealize Automation](#) 148
- [Designing and Publishing Blueprints](#) 149

About This Book

The *VMware Integrated OpenStack Administrator Guide* shows you how to perform VMware Integrated OpenStack cloud administrative tasks in the VMware Integrated OpenStack, including how to create and manage projects, users accounts, flavors, images, and networks.

Intended Audience

This guide is for cloud administrators who want to create and manage resources with an OpenStack deployment that is fully integrated with VMware[®] vSphere[®]. To do so successfully, you should be familiar with the OpenStack components and functions.

VMware Technical Publications Glossary

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

Updated Information

This *Administrator Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *Administrator Guide*.

Revision	Description
01	<ul style="list-style-type: none">■ Added Chapter 8 Working with VMware Integrated OpenStack Carrier Edition.■ Added Use Tenant Virtual Data Center for Secure Multi-Tenancy and Resource Allocation to Chapter 8 Working with VMware Integrated OpenStack Carrier Edition.■ Moved Configure Passthrough Devices for Instances into Chapter 8 Working with VMware Integrated OpenStack Carrier Edition.
00	Initial release.

About VMware Integrated OpenStack

1

With VMware Integrated OpenStack, you can implement OpenStack services on your existing VMware vSphere implementation.

You deploy VMware Integrated OpenStack through the Integrated OpenStack Manager vApp in vCenter.

The Integrated OpenStack Manager provides a workflow that guides you through and completes the VMware Integrated OpenStack deployment process. With Integrated OpenStack Manager, you can specify your management and compute clusters, configure networking, and add resources. Post-deployment, you can use Integrated OpenStack Manager to add components or otherwise modify the configuration of your VMware Integrated OpenStack cloud infrastructure.

VMware Integrated OpenStack 4.x is based on the Ocata release of OpenStack.

This section includes the following topics:

- [Internationalization](#)
- [OpenStack Foundation Compliance](#)
- [VMware Integrated OpenStack System Requirements](#)
- [OpenStack Instances in vSphere Web Client](#)
- [Monitor OpenStack Instances in the vSphere Web Client](#)
- [Customer Experience Improvement Program](#)

Internationalization

VMware Integrated OpenStack 2.0 and greater is available in English and seven additional languages: Simplified Chinese, Traditional Chinese, Japanese, Korean, French, German, and Spanish.

ASCII characters must be used for all input and naming conventions of OpenStack resources (such as project names, usernames, image names, and so on) and for the underlying infrastructure components (such as ESXi hostnames, vSwitch port group names, data center names, datastore names, and so on).

OpenStack Foundation Compliance

Every new version of VMware Integrated OpenStack complies with the latest Guidelines created by the OpenStack Foundation DefCore Committee.

VMware Integrated OpenStack is designated as an OpenStack Powered Platform™ product and therefore provides proven interoperability with all other OpenStack Powered™ products.

For detailed information about the compatibility of VMware Integrated OpenStack with the OpenStack Powered Platform™, go to

<http://www.openstack.org/marketplace/distros/distribution/vmware/vmware-integrated-openstack>.

VMware Integrated OpenStack System Requirements

Before you begin the VMware Integrated OpenStack deployment tasks, your system must comply with all hardware, software, networking, and storage requirements.

Hardware Requirements for VMware Integrated OpenStack

The hardware requirements are based on the number of VMs used for each component. For example, two VMs are used for load balancing, each of which requires two CPUs for a total requirement of four CPUs. The requirements vary depending on whether your OpenStack deployment uses Virtual Distributed Switch (VDS) or VMware NSX for vSphere (NSX) with the Networking component.

Core VMware Integrated OpenStack Components

Component	VMs	CPU	RAM (GB)	Disk Space (GB)
Integrated OpenStack Manager	1	2 (2 per VM)	4 (4 per VM)	25
Load balancing service	2	4 (2 per VM)	8 (4 per VM)	40 (20 per VM)
Database service	3	12 (4 per VM)	48 (16 per VM)	240 (80 per VM)
Memory cache service	2	4 (2 per VM)	32 (16 per VM)	40 (20 per VM)
Message queue service	2	8 (4 per VM)	32 (16 per VM)	40 (20 per VM)
Controllers	2	16 (8 per VM)	32 (16 per VM)	160 (80 per VM)
Compute service (Nova CPU)	1	2 (2 per VM)	4 (4 per VM)	20 (20 per VM)
DHCP service (VDS deployments only)	2	8 (4 per VM)	32 (16 per VM)	40 (20 per VM)
TOTAL	15	56	192	605

VMware Integrated OpenStack Requirements for Compact Mode

VMware Integrated OpenStack 3.0 supports a new deployment mode -- compact mode, which runs using a minimal amount of hardware. For information on the hardware requirements for running in compact mode, see the *VMware Integrated OpenStack Installation and Configuration Guide*.

NSX Components

Additional CPU, RAM, and disk space is required for NSX components if they are deployed with VMware Integrated OpenStack.

Component	VMs	CPU	RAM	Disk Space
NSX Controller	3	12 (4 per VM)	12 GB (4 per VM)	60 GB (20 per VM)
NSX Manager	1	4 (4 per VM)	12 GB (12 per VM)	60 GB (60 per VM)
NSX Edge (see note below)	Varies: created on demand.	1 per Edge DHCP VM, 2 per Edge router VM	512 MB per Edge DHCP VM, 1 per Edge router VM	512 MB per Edge DHCP VM, 1 per Edge router VM
TOTAL	4 plus Edge requirements	16 plus Edge requirements	24 GB plus Edge requirements	120 GB plus Edge requirements

When you create a logical subnet or logical router, a new Edge VM is dynamically created to serve this request if an existing Edge node cannot.

Software Requirements for VMware Integrated OpenStack

Before you begin the VMware Integrated OpenStack deployment tasks, the software components must meet all of the version prerequisites for vSphere, ESXi hosts, and the NSX product.

Requirement	Description
vSphere version	<ul style="list-style-type: none"> ■ vSphere 5.5 Update 2 Enterprise Plus ■ vSphere 6 Enterprise Plus
ESXi hosts	<ul style="list-style-type: none"> ■ Version 5.5 Update 2 ■ Eight or more logical processors on each host. ■ The vCenter and all ESXi hosts intended for the VMware Integrated OpenStack deployment must use the same Network Time Protocol (NTP) server. ■ Verify that the ESXi host firewalls are configured to allow gdbserver access. Typically, the port range is 5900-5964.
NSX	Consult with VMware for the preferred version.

Storage Requirements for NSX Deployments

Storage requirements vary depending on your deployment configuration. Different nodes and clusters can share datastores. For example, during the installation process, you can specify the same datastore for the Compute and Image Service nodes.

For information about storage requirements per VM in a typical VMware Integrated OpenStack deployment, see [Hardware Requirements for VMware Integrated OpenStack](#).

Storage requirements vary depending on whether you deploy with NSX or VDS networking.

Storage Requirements for NSX Deployments

NSX Controller, Manager, and Edge nodes affect the storage needs in an NSX deployment.

Cluster	Storage Requirements (GB)	Notes
Management	665	<p>The storage requirement calculation is based on the following nodes:</p> <ul style="list-style-type: none"> ■ OpenStack Manager (1 node) ■ Load Balancers (2 nodes) ■ Database (3 nodes) ■ Memory Cache (2 nodes) ■ Message Queue (2 nodes) ■ Controllers (2 nodes) ■ NSX Controller (3 nodes) ■ NSX Manager (1 node)
Compute	20	<p>Value is per cluster.</p> <p>Each Compute cluster contains a single Compute node. To add capacity, add clusters.</p>
NSX Edge	1.5	<p>Value is per node.</p> <p>Storage requirements for the NSX Edge cluster vary. When you create a logical subnet or router but an existing NSX Edge node cannot serve the request, an additional node is dynamically created.</p> <p>Note Creating a dedicated cluster for the NSX Edge nodes is a best practice to optimize performance. In an alternative deployment, you can include the NSX Edge nodes in the Management cluster.</p>

Storage Requirements for VDS Deployments

DHCP nodes affect the storage needs in a VDS deployment.

Cluster	Storage Requirements (GB)	Notes
Management	585	The storage requirement calculation is based on the following service nodes: <ul style="list-style-type: none"> ■ OpenStack Manager (1 node) ■ Load Balancers (2 nodes) ■ Database (3 nodes) ■ Memory Cache (2 nodes) ■ Message Queue (2 nodes) ■ Controllers (2 nodes) ■ DHCP Controller (2 nodes)
Compute	20	Value is per cluster. Each Compute cluster contains a single Compute node. To add capacity, add clusters.

Required NSX Parameters

When you are deploying VMware Integrated OpenStack with NSX for the Networking component, you must configure the NSX nodes in advance.

When you install VMware Integrated OpenStack, you must provide the following information.

Starting with VMware Integrated OpenStack 3.1, if you use VMware NSX-T in your environment, you can use the native DHCP and metadata support. To be able to use these functionalities, you must create a DHCP profile and metadata proxy server for your NSX-T environment.

Property	Description
Username	User name for accessing the NSX Manager node.
Password	Password for accessing the NSX Manager node.
Transport Zone	Name of the default transport zone.
Edge Cluster	The name of the cluster containing the Edge nodes.
Virtual Distributed Switch for Edge VTEP	The VDS from the NSX configuration.
Port Group for External Network	The port group created on a VLAN specifically for the External network. You created this port group as part of the process of preparing to deploy VMware Integrated OpenStack with NSX.
(optional VMware NSX-T only) DHCP profile	To use native DHCP, configure a DHCP server profile for your NSX-T environment. For more information, see <i>Create a DHCP Server Profile</i> in the <i>NSX-T Administration Guide</i> .
(optional VMware NSX-T only) Metadata proxy server	To use metadata support, configure a metadata proxy server for your NSX-T environment. For more information, see <i>Add a Metadata Proxy Server</i> in the <i>NSX-T Administration Guide</i> . During the configuration, use the load balancer private IP of your OpenStack deployment for URL for the Nova server. For example: <code>http://load_balancer_private_IP:8775/</code> . Also keep the secret parameter, as you need it during the VMware Integrated OpenStack deployment.

OpenStack Instances in vSphere Web Client

The VMs you create in your VMware Integrated OpenStack deployment appear in your vCenter inventory. Many restrictions apply to how you manage and work with OpenStack VMs.

In most cases, you must manage such VMs in the OpenStack dashboard or CLI rather than in the vSphere Web Client.

OpenStack Features Supported in vSphere

vSphere supports certain OpenStack features.

OpenStack Feature	Supported in vSphere
Launch	YES
Reboot	YES
Terminate	YES
Resize	YES
Rescue	YES
Pause	NO
Un-pause	NO
Suspend	YES
Resume	YES
Inject Networking	
Inject Networking is supported only when the following conditions are present:	
<ul style="list-style-type: none"> ■ With nova network in Flat mode ■ With Debian- or Ubuntu-based virtual machines ■ At boot time 	YES
Inject File	NO
Serial Console Output	YES
RDP Console	NO
Attach Volume	YES
Detach Volume	YES
Live Migration	YES
Snapshot	YES
iSCSI	YES
Fibre Channel	YES
	Supported through vSphere datastores
Set Admin Pass	NO
Get Guest Info	YES

OpenStack Feature	Supported in vSphere
Set Host Info	YES
Glance Integration	YES
Service Control	YES
VLAN Networking	YES
Flat Networking	YES
Security Groups	NO vSphere Web Client supports Security Groups when using the Neutron plugin of VMware NSX for vSphere .
Firewall Rules	NO
Routing	YES
Config Drive	YES
Evacuate or Host Maintenance Mode	YES
Volume Swap	NO
Volume Rate Limiting	NO

VM Operations in OpenStack

The following table maps VMware Integrated OpenStack and vSphere VM operations, and provides recommendations about where best to perform the operation. If you create a VM in VMware Integrated OpenStack, manage that VM in VMware Integrated OpenStack.

vSphere Feature	OpenStack Counterpart	Exposed through OpenStack API	Where to Perform this Operation
Create a virtual machine	Launch instance	YES	OpenStack dashboard
Reboot	Reboot	YES	OpenStack dashboard or vSphere Web Client
Delete	Terminate	YES	OpenStack dashboard
Resize	Resize	YES	OpenStack dashboard
Pause	Pause	YES	OpenStack dashboard or vSphere Web Client
Unpause	Un-pause	YES	OpenStack or vSphere Web Client
Pause	Suspend	YES	OpenStack dashboard
Resume	Resume	YES	OpenStack dashboard
Serial Console Output	Serial Console Output	YES	OpenStack dashboard or vSphere Web Client
RDP Console	RDP Console		OpenStack dashboard or vSphere Web Client
Add Disk	Attach Volume	YES	OpenStack dashboard
Remove Disk	Detach Volume	YES	OpenStack dashboard

vSphere Feature	OpenStack Counterpart	Exposed through OpenStack API	Where to Perform this Operation
vMotion	Live Migration	YES	vSphere Web Client
Snapshot	Snapshot	YES	OpenStack dashboard or vSphere Web Client
Functions available through VMware Tools .	Get Guest Info/Get Host Info	YES	OpenStack dashboard or vSphere Web Client For vSphere Web Client, this function is available with VMware Tools.
Distributed Port Groups	VLAN Networking or Flat Networking	YES	OpenStack dashboard
Function available through VMware Tools.	Config Drive	NO	OpenStack dashboard or vSphere Web Client For vSphere Web Client, this function is available with VMware Tools.
InstallVMware Tools in a VM	Install VMware Tools in a VM	NO	OpenStack dashboard or vSphere Web Client

vCenter Features Not Supported in the OpenStack API

Direct parity does not exist between OpenStack features and vSphere features. The OpenStack API does not support the following vCenter features.

- Adding a host to a cluster

OpenStack cannot add a host to a cluster in vSphere.

- Putting a host into maintenance mode

You place a host in maintenance mode to service it, for example, to install more memory. A host enters or leaves maintenance mode only as the result of a user request. No such function exists in OpenStack. See the vSphere documentation for instructions about entering and exiting maintenance mode.

- Resource Pools

A resource pool in vSphere is a logical abstraction for flexible management of resources, such as CPU and memory. OpenStack has no equivalent to a resource pool.

- vSphere snapshots

vCenter supports OpenStack snapshots, but vSphere snapshots are distinct and are not supported in the OpenStack API.

Monitor OpenStack Instances in the vSphere Web Client

You can view and monitor instance activity and metadata in the vSphere Web Client.

Prerequisites

Verify that VMware Integrated OpenStack is deployed and operational.

Verify that you or another user has started instances in VMware Integrated OpenStack.

Procedure

1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.

2 Expand the Inventory view until you expose the instance VMs in the compute cluster.

The instance VMs are identified by their UUIDs.

3 Select an instance VM and click the **Summary** tab.

The **Summary** tab displays the portlets common to VMs in the vSphere Web Client. The OpenStack VM and Tags portlets contain details about instances created in OpenStack.

4 (Optional) Locate and review the OpenStack VM and Tags VM portlets.

These portlets display information about the selected instance, including instance properties such as the name, tenant, the user that created the instance, the originating flavor, and so on.

5 (Optional) Use the vSphere Web Client to search for and filter OpenStack instances.

a In the vSphere Web Client Search field, enter one of the tag values from the Tags portlet.

For example, to find all instances created using the default m1.tiny flavor, enter **m1.tiny**.

The **Related Objects** tab appears with a list of all the OpenStack instances that match the search criteria.

b Click the name of any instance to open the **Summary** tab for that instance.

Customer Experience Improvement Program

You can configure this product to collect data that can be used by the VMware Customer Experience Improvement Program.

This product participates in VMware's Customer Experience Improvement Program ("CEIP"). Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>

To join or leave the CEIP for this product, please go to the Customer Experience Improvement Program page in the User Interface to change your participation in CEIP:

- During product deployment using the Integrated OpenStack Manager, participation in the CEIP is enabled by default, unless you choose not to participate.

After initial deployment, go to the Customer Experience Improvement Program page to modify your participation, if required.

- To join the CEIP, go to **Home > Inventories**, and click the VMware Integrated OpenStack icon. Then click the **Manage** tab and click the **Settings** tab. Finally click **Enable** to join.
- To leave the CEIP, go to **Home > Inventories**, and click the VMware Integrated OpenStack icon. Then click the **Manage** tab and click the **Settings** tab. Finally click **Disable** to leave the program.

Managing Your VMware Integrated OpenStack Deployment

2

Managing your VMware Integrated OpenStack deployment includes modifying configuration settings, backup, recovery, and restoration of your OpenStack configuration and data; using patches for minor updates, and upgrading to new versions.

This section includes the following topics:

- [Managing Your Deployment Configuration](#)
- [Managing Your Network Configuration](#)
- [Adding Capacity in the vSphere Web Client](#)
- [Configure the Backup Service for Block Storage](#)
- [Back Up the VMware Integrated OpenStack Deployment](#)
- [Restore VMware Integrated OpenStack from a Backup](#)
- [Failure Recovery](#)
- [VMware Integrated OpenStack Log File Locations](#)
- [Upgrade to VMware Integrated OpenStack 4.0](#)
- [Updating Your VMware Integrated OpenStack Deployment](#)
- [Customize the Dashboard Logos and Background](#)
- [Use Profiling to Trace OpenStack Deployments](#)
- [Configure NUMA for Use With VMware Integrated OpenStack](#)

Managing Your Deployment Configuration

During the VMware Integrated OpenStack installation and deployment process, you configure the OpenStack components, specify the syslog server, provide passwords for LDAP, NSX, and vCenter Server, among other deployment tasks. After deployment, you can modify these settings.

Monitor Your VMware Integrated OpenStack Deployment

After you finish installing VMware Integrated OpenStack, you can monitor your deployment configuration, including datastore sizes, network settings, and metadata service, among other details.

Procedure

- 1 In vCenter, select **Home > VMware Integrated OpenStack**.
- 2 Click the **Monitor** tab.

Modify the Syslog Server Address

The Syslog server address is configured during installation but you can modify the configuration afterward.

Prerequisites

Verify that the new Syslog server address is valid.

Procedure

- 1 In vCenter, select **Home > VMware Integrated OpenStack > Manage**.
- 2 Click the **Settings** tab.
- 3 Click **Syslog Server**.

The Syslog Server panel displays the current configuration.

- 4 Click **Edit** to change the Syslog server address.
- 5 Click **OK** to apply the change.

The vSphere Web Client might take a few minutes to update the OpenStack configuration.

Update Deployment Passwords

Part of your VMware Integrated OpenStack deployment configuration includes passwords that allow OpenStack to access and connect with your LDAP server, NSX, and vCenter Server. If the credentials change, you can modify the password settings directly in the VMware Integrated OpenStack manager to ensure continued access.

Only the text boxes with updated passwords on the Change Password page change. To leave a password unmodified, leave the text box blank.

Prerequisites

Verify that the passwords you supply in the Change Passwords panel match the passwords configured for the LDAP server, NSX, or vCenter Server, as appropriate.

Procedure

- 1 In vCenter, select **Home > VMware Integrated OpenStack > Manage**.

- 2 Click the **Settings** tab.
- 3 Click **Change Password**.

The Change Passwords panel contains text boxes for updating the current LDAP server, NSX, and vCenter Server password configurations.

- 4 Enter the new password.
- 5 Click **Submit**.

The password settings in the VMware Integrated OpenStack are updated with the new values.

Manage the OpenStack SSL Certificate Configuration

You can add OpenStack SSL certificates in the VMware Integrated OpenStack manager.

You can only import existing CA signed certificates, created from CSRs generated by VMware Integrated OpenStack. You can also create new CSRs to create new CA signed certificates. Using wildcard certificates is not supported.

Procedure

- 1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.
- 2 Click the **Manage** tab and click the **Settings** tab.
- 3 Click **OpenStack SSL Certificate**.
- 4 Generate a new certificate signing request to create new CA signed certificate.
 - a Provide the Organizational Unit, Organizational Name, Locality Name, State Name, and Country Code information as appropriate to your organization.
 - b Click **Generate**.
 - c Use the generated certificate signing request to create a certificate that is signed by your CA.
- 5 Import the CA signed certificate.
 - a Click **Import**.
 - b Browse to and select the CA signed certificate file.
 - c Click **OK**.

The imported certificate is applied.

Configure the Ceilometer Component

Ceilometer is the telemetric component of OpenStack that collects and persists data regarding the use of the physical and virtual resources in your OpenStack deployment.

You can enable Ceilometer after completing the VMware Integrated OpenStack deployment.

Procedure

1 In vCenter, select **Home > VMware Integrated OpenStack > Manage**.

2 Select the **Settings** tab.

3 Click **Ceilometer**.

The Ceilometer panel displays the current status and configuration.

4 Click **Edit** to modify the settings.

5 Select the **Configure Ceilometer** option.

6 Click **OK** to configure Ceilometer.

The vSphere Web Client might take a few minutes to update the OpenStack configuration.

Ceilometer is automatically enabled the first time you configure it. Afterwards, the Ceilometer settings show only **Enable** and **Disable** options.

Modify Your Enrollment in the Customer Experience Improvement Program

During the installation process, you can enroll in the VMware Customer Experience Improvement Program (CEIP). After installation, you can modify this configuration in the VMware Integrated OpenStack manager.

This product participates in VMware's Customer Experience Improvement Program ("CEIP"). Details regarding the data collected through CEIP and the purposes for which it is used by VMware are set forth at the Trust & Assurance Center at <http://www.vmware.com/trustvmware/ceip.html>

Procedure

1 To join the CEIP, go to **Home > Inventories**, and click the VMware Integrated OpenStack icon. Then click the **Manage** tab and click the **Settings** tab. Finally, click **Enable** to join.

2 To leave the CEIP, go to **Home > Inventories**, and click the VMware Integrated OpenStack icon. Then click the **Manage** tab and click the **Settings** tab. Finally, click **Disable** to leave.

3 Click **Customer Experience Improvement Program**.

The Customer Experience Improvement Program page displays the current status of your participation in the CEIP. If enabled, you are opted in. If disabled, you are opted out.

Manage Your Authentication Settings

Part of your VMware Integrated OpenStack deployment configuration includes setting up authentication. You can also modify this configuration post-installation.

Prerequisites

Verify that the new LDAP settings are valid.

Procedure

1 In vCenter, select **Home > VMware Integrated OpenStack > Manage**.

2 Click the **Settings** tab.

3 Click **Configure Identity Source**.

The panel displays the current configuration.

4 Set the VMware Integrated OpenStack identity source.

Option	Description
OpenStack admin user	Define the OpenStack administrative user name. This is the default administrative user name for logging in to the VMware Integrated OpenStack dashboard.
OpenStack admin password	Define the OpenStack administrative user password. This is the default administrative user password for logging in to the VMware Integrated OpenStack dashboard.
Confirm password	Reenter the password for confirmation.

5 If you are using LDAP with your VMware Integrated OpenStack deployment, click the plus sign (+) to configure the LDAP source.

The Add Identity Source dialog appears.

Option	Description
Domain Name	Specify the full Active Directory domain name; for example, vmware.com.
Bind user	Provide the user name to bind to Active Directory for LDAP requests.
Bind password	Provide the password to allow the LDAP client access to the LDAP server.
Domain controllers	(Optional) VMware Integrated OpenStack automatically chooses the existing Active Directory domain controllers. However, you can specify a list of specific domain controllers to use. To do this, select the Domain controllers radio button and then enter the IP address of one or more domain controllers, separated by commas.
Site	(Optional) Optionally, you can limit LDAP searching to a specific deployment site within your organization; for example, sales.vmware.com. Select the Site radio button and enter the domain name of the site to search.
User Tree DN	(Optional) Enter the search base for users; for example, DC=vmware, DC=com. Defaults to the top of the user tree in most Active Directory deployments.
User Filter	(Optional) Enter an LDAP search filter for users. Important If you use VMware Integrated OpenStack 3.0 or older and your directory contains more than 1,000 objects (users and groups), you must apply a filter to ensure that fewer than 1,000 objects are returned. For examples of filters, see https://msdn.microsoft.com/en-us/library/aa746475(v=vs.85).aspx .
Advanced setting	If you want to specify advanced LDAP settings, check the Advanced setting check box.

If you check the **Advanced setting** check box, additional LDAP configuration fields appear.

Note Always contact the LDAP administrator to obtain correct values for advanced LDAP settings, or use tools such as `ldapsearch` or `Apache Directory Studio` to locate the settings.

Option	Description
Encryption	From the pull-down menu, choose None , SSL , or StartTLS
Hostname	Enter the hostname for the LDAP server.
Port	Enter the port number to user on the LDAP server.
User objectclass	(Optional) Enter the LDAP object class for users.
User ID attribute	(Optional) Enter the LDAP attribute mapped to the user ID. Note that this value cannot be a multi-valued attribute.
User name attribute	(Optional) Enter the LDAP attribute mapped to the user name.
User mail attribute	(Optional) Enter the LDAP attribute mapped to the user email.
User password attribute	(Optional) Enter the LDAP attribute mapped to the password.
Group objectclass	(Optional) Enter an LDAP object class for groups.
Group ID attribute	(Optional) Enter the LDAP attribute mapped to the group ID.
Group name attribute	(Optional) Enter the LDAP attribute mapped to the group name.
Group member attribute	(Optional) Enter the LDAP attribute mapped to the group member name.
Group description attribute	(Optional) Enter the LDAP attribute mapped to the group description.

Figure 2-1. Add identity source dialog

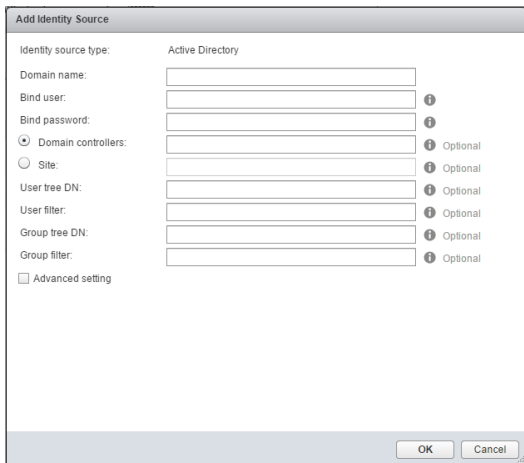


Figure 2-2. Advanced LDAP settings

<input checked="" type="checkbox"/> Advanced setting		
Encryption:	SSL	
Host name:		
Port:	389	
User objectclass:		Optional
User ID attribute:		Optional
User name attribute:		Optional
User mail attribute:		Optional
User password attribute:		Optional
Group objectclass:		Optional
Group ID attribute:		Optional
Group name attribute:		Optional
Group member attribute:		Optional
Group description attribute:		Optional

6 Click **Save**.

What to do next

To complete the LDAP configuration, you must manually modify the default OpenStack domain configuration. See [Modify the Default Domain Configuration](#).

Modify the Default Domain Configuration

By default, the Identity Service component (Keystone) does not return users and groups to the default domain. The following procedure modifies the default configuration to ensure that users with administrative privileges can access and assign LDAP users to roles in OpenStack.

Prerequisites

- Verify that you have successfully deployed VMware Integrated OpenStack.
- Verify that VMware Integrated OpenStack is running.
- Verify that Active Directory is configured as the LDAP backend.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack deployment.

This step varies depending on your mode of deployment.

- If your deployment is using compact mode, log into the controller node.
- If your deployment is high-availability mode, log into the load balancer node.

- 2 Switch to root user.

```
sudo su -
```

- 3 Execute the `cloudadmin_v3.rc` file.

```
$ source ~/cloudadmin_v3.rc
```


4 Create the initial project in the default domain in OpenStack.

```
$ openstack --os-identity-api-version 3 --os-username admin \
  --os-user-domain-name local --os-project-name admin --os-password admin \
  --os-region-name nova project create --domain default --description "Demo Project" --or-
  show demo
```

Parameter	Description
<code>--os-identity-api-version 3</code>	Specifies the API version, in this case, version 3.
<code>--os-username admin</code>	Provides the administrative username for login, in this case admin .
<code>--os-user-domain-name local</code>	Specifies the domain, in this case local for the specified user.
<code>--os-project-name admin</code>	Specifies the admin OpenStack project.
<code>--os-password admin</code>	Provides the administrative password for login, in this case admin .
<code>--os-region-name nova project create</code>	Runs the <code>nova project create</code> command.
<code>--domain default</code>	This command specifies the domain where the new project is created, in this case the default domain.
<code>--description "Demo Project"</code>	This parameter names the new project, in this case Demo Project .
<code>--or-show demo</code>	Creates an alias for the new project.

5 Add an administrative user to the new project in the default domain.

```
$ openstack --os-identity-api-version 3 --os-username admin \
  --os-user-domain-name local --os-project-name admin --os-password admin \
  --os-region-name nova role add --project demo --project-domain default \
  --user SOMEUSER@vmware.com --user-domain default admin
```

Parameter	Description
<code>--os-identity-api-version 3</code>	Specifies the API version, in this case, version 3.
<code>--os-username admin</code>	Provides the administrative username for login, in this case admin .
<code>--os-user-domain-name local</code>	Specifies the domain, in this case local for the specified user.
<code>--os-project-name admin</code>	Specifies the admin OpenStack project.
<code>--os-password admin</code>	Provides the administrative password for login, in this case admin .
<code>--os-region-name nova role add</code>	Runs the <code>nova role add</code> command.
<code>--project demo</code>	Specifies the project to which the new administrative user is added.
<code>--project-domain default</code>	Specifies the project domain.

Parameter	Description
<code>--user SOMEUSER@vmware.com</code>	Specifies the new administrative user.
<code>--user-domain default admin</code>	Assigns the new user to the default admin domain.

Note If special characters are used for the user ID, you must modify the Keystone settings in the VMware Integrated OpenStack manager.

- 6 (Optional) If special characters are used for the administrative user ID, you must modify the Keystone settings in the VMware Integrated OpenStack manager.
 - a In the VMware Integrated OpenStack manager in vCenter, go to **Manage > Settings > Configure Identity Source**.
 - b Click **Edit**.
 - c Under Advanced Settings, modify the User ID value from `cn` to `userPrincipalName`.

You can now log in to the default domain in the VMware Integrated OpenStack dashboard using the administrative user name and password.

Configure VMware Identity Manager as a Single Sign-On Solution for OpenStack

Starting with VMware Integrated OpenStack 3.1, you can integrate your VMware Integrated OpenStack deployments with VMware Identity Manager.

By integrating VMware Integrated OpenStack with VMware Identity Manager you achieve a way to securely use existing credentials to access cloud resources such as servers, volumes, and databases, across multiple endpoints provided in multiple authorized clouds. You have a single set of credentials, without having to provision additional identities or log in multiple times. The credential is maintained by the user's Identity Provider.

Prerequisites

- Verify that the version of VMware Identity Manager is 2.8.0 or later.
- Verify that you can authenticate as administrator to the VMware Identity Manager instance.

Procedure

- 1 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Edit the `/opt/vmware/vio/custom/custom.yml` file in a text editor to configure it for your environment.
 - a Under **Federation**, uncomment the following parameters and set values for your environment.

The following example provides guidance for the most common configuration with VMware Identity Manager.

Parameter	Value
<code>federation_protocol</code>	saml2
<code>federation_idp_id</code>	vidm
<code>federation_idp_name</code>	VIDM SSO
<code>federation_idp_metadata_url</code>	<code>https://IDP_HOSTNAME/SAAS/API/1.0/GET/metadata/idp.xml</code>
<code>federation_group</code>	Federated Users
<code>federation_group_description</code>	Groups for all federated users
<code>vidm_address</code>	<code>IDP_URL</code>
<code>vidm_user</code>	<code>vidm_administrative_user</code>
<code>vidm_password</code>	<code>vidm_administrative_user_password</code>
<code>vidm_insecure</code>	False
<code>vidm_group</code>	ALL USERS

- b Save the `custom.yml` file.
- 3 Enable federation with the settings that you configured in the `custom.yml` file.

```
viocli deployment configure --tags federation --limit controller,lb
```

After the integration operation completes successfully, the VMware Integrated OpenStack dashboard shows a new **Authenticate using** drop-down menu that allows the user choose the authentication method.

- 4 Prior to being able to login a VMware Identity Manager user to VMware Integrated OpenStack, assign a role/project to the group that user belongs to.

You might have to create a group in keystone that corresponds to a group found in VMware Identity Manager that a user is a member of. For VMware Identity Manager users, Keystone does not automatically create groups but ephemeral users. If the group does not exist, the user becomes a member of the default `Federated Users` group.

- a Log in to the VMware Integrated OpenStack dashboard as an administrator.
- b Under **Federation**, click **Mappings** to see the current mappings.
- c Click **Edit** to configure a mapping according to your needs.

For more information about mappings, see the [Mapping Combinations for Federation](#) in the OpenStack documentation.

Managing Your Network Configuration

During installation, you configure the Neutron networking component by specifying port groups. After installation, you can extend the IP range, create an L2 bridge, or change the DNS of the dedicated networks.

Add IP Address Ranges to a Network

You can add IP address ranges to the management or API access network.

You typically add IP ranges as part of the upgrade process.

Procedure

- 1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.

- 2 Click the **Manage** tab and click the **Networks** tab.

The **Networks** tab lists the Management and API network configurations, including their IP address ranges.

- 3 Expand the IP addresses available for the Management network.

- a Right-click the name of the Management network in the list and select **Add IP Range**.
- b In the Add IP Range dialog box, specify the new IP range.

Note If you are adding addresses as part of the upgrade process, the new IP range must match the same number of IP addresses configured for the existing Management network. For example, in a typical VMware Integrated OpenStack deployment, the Management network requires a minimum range of 11 IP addresses.

- c Click **OK**.

- 4 Expand the IP addresses available for the external network.

- a Right-click the name of the API network in the list and select **Add IP Range**.
- b In the Add IP Range dialog box, specify the new IP range.

Note If you are adding addresses as part of the upgrade process, the new IP range must match the same number of IP addresses configured for the existing API network. For example, in a typical VMware Integrated OpenStack deployment, the API network requires a minimum range of 2 IP addresses for versions 3.0 and older and a minimum range of 3 IP addresses for versions 3.1 and later.

- c Click **OK**.

Modify the Default Router Setting

You can modify the default router setting that NSX uses in the `custom.yml` file.

The Neutron configuration file includes a parameter that sets the default router types. For example, `tenant_router_types = shared, distributed, exclusive`. You can modify the `custom.yml` file to override this configuration with a custom setting.

Procedure

- 1 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
- 3 Uncomment the `nsxv_tenant_router_types` parameter and specify the router types for NSX tenants.

```
nsxv_tenant_router_types: exclusive, shared, distributed
```

- 4 Using SSH, log in to the VMware Integrated OpenStack manager.
- 5 Switch to root user.

```
sudo su -
```

- 6 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment configure --limit controller
```

Modify Network DNS Setting

After installation, you can modify the DNS settings for the networks configured for OpenStack management and API access.

Important Modifying the network DNS setting results in a brief interruption in the network connection.

Procedure

- 1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.
- 2 Click the **Manage** tab and click the **Networks** tab.

The **Networks** tab lists the Management and API network configurations, including their DNS addresses.

- 3 Right-click the network name whose DNS setting you want to modify and choose **Change DNS**.

Note You can also select the network in the list, click **All Actions** and choose **Change DNS**.

- 4 Modify the DNS and Secondary DNS IP addresses.
- 5 Click **OK**.

Create a VXLAN/VLAN L2 Bridge

In a leaf-spine data center architecture, the OpenStack Compute cluster cannot access VMs on a VLAN. You can overcome this technical limitation by creating a VXLAN network and L2 VXLAN and VLAN bridge.

Prerequisites

Verify that a VDS port group is available for the VXLAN network configuration.

Procedure

- 1 Using SSH, log in as administrator to the VMware Integrated OpenStack manager.
- 2 Using SSH, log in to the controller01 node.
- 3 Create the logical L2 gateway on Neutron.
 - ◆ If your version is VMware Integrated OpenStack 3.0 or older, use the `neutron-l2gw l2-gateway-create` command.

```
neutron-l2gw l2-gateway-create <gateway-name> \
--device name=<device-name1>,interface_names="<interface-name1>[|<seg-id1>]"
```

- ◆ If your version is VMware Integrated OpenStack 3.1 or later, use the `l2-gateway-create` command.

```
l2-gateway-create <gateway-name> \
--device name=<device-name1>,interface_names="<interface-name1>[|<seg-id1>]"
```

Option	Description
<gateway-name>	Specifies the name of the new gateway.
<device-name1>	Specifies the device name. This is a dummy name. The NSX plug-in creates a dedicated DLR.
<interface-name1>	Specifies the distributed port group MOB ID as the interface name.
<seg-id1>	Specifies the distributed port group segmentation ID.

From the backup edge pool, NSX creates a dedicated DLR called L2 bridging-{gateway-id}.

4 Create the logical L2 gateway connection on Neutron.

- ◆ If your version is VMware Integrated OpenStack 3.0 or older, use the `neutron-l2gw l2-gateway-connection-create` command.

```
neutron-l2gw l2-gateway-connection-create <gateway-name/uuid> <network-name/uuid> \
[--default-segmentation-id=<seg-id>]
```

- ◆ If your version is VMware Integrated OpenStack 3.1 or later, use the `l2-gateway-connection-create` command.

```
l2-gateway-connection-create <gateway-name/uuid> <network-name/uuid> \
[--default-segmentation-id=<seg-id>]
```

Option	Description
<code><gateway-name/uuid></code>	Specifies the name of the existing gateway.
<code><network-name/uuid></code>	Specifies the network name. This is a dummy name. The NSX plug-in creates a dedicated DLR.
<code><default-segmentation-id=seg-id1></code>	Specifies the default distributed port group segmentation ID.

This operation connects the OpenStack network with the Provider VLAN network.

Managing NSX Edge Node High Availability

You can configure VMware Integrated OpenStack to ensure that every NSX Edge node is enable for high availability.

You can configure the `custom.yml` file before installing and deploying VMware Integrated OpenStack. If you have already installed and deployed VMware Integrated OpenStack, you have the additional step of manually enabling each running NSX Edge node.

Enable NSX Edge Node High Availability Before Deployment

Before you install VMware Integrated OpenStack, you can enable high availability for NSX Edge nodes by modifying the `custom.yml` file.

Prerequisites

Verify that your Edge cluster has at least two hosts. If not, you might receive an anti-affinity error.

Procedure

- 1 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Edit the `/opt/vmware/vio/custom/custom.yml` file.
 - a Uncomment the `nsxv_edge_ha` parameter.
 - b Set the `nsxv_edge_ha` parameter to **True**.

```
nsxv_edge_ha: True
```

- 3 Save the `custom.yml` file.

When you install and deploy VMware Integrated OpenStack, high availability is enabled by default for all NSX Edge nodes.

Enable NSX Edge Node High Availability After Deployment

If you have already installed VMware Integrated OpenStack, you can enable high availability for NSX Edge nodes by modifying the `custom.yml` file and manually modifying each running Edge node.

Prerequisites

Verify that your Edge cluster has at least two hosts. If not, you might receive an anti-affinity error.

Procedure

- 1 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Edit the `/opt/vmware/vio/custom/custom.yml` file.
 - a Uncomment the `nsxv_edge_ha` parameter.
 - b Set the `nsxv_edge_ha` parameter to **True**.

```
nsxv_edge_ha: True
```

- 3 Save the `custom.yml` file.

After modifying and saving the `custom.yml` file, high availability is enabled for newly deployed NSX Edge nodes subsequently generated by VMware Integrated OpenStack.

- 4 Manually enable high availability on all current NSX Edge nodes.
 - a In the VMware Integrated OpenStack controller, get a list of all current Edge nodes and their edge-id values.

```
sudo -u neutron nsxadmin -r edges -o nsx-list
```

- b Enable high availability on each Edge node by specifying its edge-id value.

```
sudo -u neutron nsxadmin -r edges -o nsx-update \  
--property highAvailability=True \  
--property edge-id=<edge-id>
```

- c Repeat the preceding command for each Edge node.

- 5 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment -v configure
```

Important This command updates your entire deployment and might briefly interrupt operations.

Disable Public API Access

You can temporarily block user access to your VMware Integrated OpenStack deployment. For example, you might need to perform maintenance tasks that require blocking users while still allowing administrative access.

By modifying the `custom.yml` file, you can block user access through the public API network. When users attempt to access OpenStack, they will see maintenance web page instead.

Procedure

- 1 If you have not already done so, implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom  
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Edit the `custom.yml` file by uncommenting the `haproxy_custom_maintenance_page` parameter.

```
#####  
# haproxy maintenance page  
#####  
# location of the maintenance page to be displayed when the public VIP is disabled  
haproxy_custom_maintenance_page : "/home/viouser/custom/503maintenance.html"  
# mail contact for maintenance page.  
#haproxy_mailto: test@vmware.com
```

- 3 Save the `custom.yml` file.

- 4 Push the modified configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment -v configure --limit lb
```

- 5 To remove the block, repeat the procedure and re-comment out the `haproxy_custom_maintenance_page` parameter.

Configure BGP Dynamic Routing for Your VMware Integrated OpenStack Deployment

Starting with VMware Integrated OpenStack 4.0, you can configure dynamic routing for your provider and tenants.

You must first create a VXLAN external network that you later use as internal interface for your gateway edges.

Prerequisites

- You must use VMware NSX for vSphere as your virtual network provider.

Procedure

- 1 Create IPv4 address scope for future tenant subnets and the external VXLAN network subnet.

```
neutron address-scope-create scope_name 4
```

- 2 Create a provider subnet pool.

Replace *scope_name* with the name of the address scope that you created earlier.

```
neutron subnetpool-create --pool-prefix 10.10.10.0/24 --default-prefixlen 24 provider_pool_name --address-scope scope_name
```

- 3 Create a self-service subnet pool for tenant networks.

Replace *scope_name* with the name of the address scope that you created earlier.

```
neutron subnetpool-create --pool-prefix 1.1.1.0/24 --default-prefixlen 26 selfservice --address-scope scope_name --shared
```

- 4 Create the external VXLAN network.

The following command creates a new logical switch in VMware NSX for vSphere .

```
neutron net-create --provider:network_type vxlan --router:external external_VXLAN_network_name
```

5 Create the external VXLAN subnet.

Replace *provider_pool_name* with the name of the provider pool that you created earlier. Replace *external_VXLAN_network_name* with the name of the network that you created earlier.

```
neutron subnet-create --no-gateway --name ext_vxlan_subnet_name --disable-dhcp --allocation-pool
start=start_IP,end=end_IP --subnetpool provider_pool_name external_VXLAN_network_name NETWORK[CIDR]
```

6 Create BGP peering gateway edges by using the nsxadmin utility.

Gateway edges use the management network as external interface and the external network that you created as internal interface.

```
nsxadmin -r bgp-gw-edge -o create --property name=name_GW-EDGE1 --property local-as=65001 --
property external-iface=morefid:mgtnetwork --property internal-
iface=morefid:internal_interface_network_GW-EDGE1
```

```
nsxadmin -r bgp-gw-edge -o create --property name=name_GW-EDGE2 --property local-as=65001 --
property external-iface=morefid:mgtnetwork --property internal-
iface=morefid:internal_interface_network_GW-EDGE2
```

7 Update the NSX Edges with BGP advertisement.

Use the IDs of the edges that you created in the previous step.

```
nsxadmin -r routing-redistribution-rule -o create --property gw-edge-ids=edge-ID_GW-EDGE1,edge-
ID_GW-EDGE2 --property learner-protocol=bgp --property learn-from=connected,bgp --property
action=permit
```

8 Update the NSX Edges with BGP neighbors.

Use the IDs of the edges that you created earlier.

```
nsxadmin -r bgp-neighbour -o create --property gw-edge-ids=edge-ID_GW-EDGE1,edge-ID_GW-EDGE2 --
property ip-address=IP_physical_router1 --property remote-as=65000 --property password=BGP_password
```

```
nsxadmin -r bgp-neighbour -o create --property gw-edge-ids=edge-ID_GW-EDGE1,edge-ID_GW-EDGE2 --
property ip-address=IP_physical_router2 --property remote-as=65000 --property password=BGP_password
```

9 Update your physical routers.

- a Set AS value to **65000**.
- b Set BGP neighbours to *name_GW-EDGE1* and *name_GW-EDGE2*.
- c Set to advertise itself as dynamic gateway.

10 Create and configure the BGP Speaker.

- a Create the BGP speaker.

```
neutron bgp-speaker-create --local-as local_as_value name_bgp_speaker
```

- b Create BGP peers.

```
neutron bgp-peer-create --peer-ip internal_interface_network_GW-EDGE1 --remote-as 65001 --password BGP_password --auth-type md5 name_GW-EDGE1 --esg-id edge-ID_GW-EDGE1
```

```
neutron bgp-peer-create --peer-ip internal_interface_network_GW-EDGE2 --remote-as 65001 --password BGP_password --auth-type md5 name_GW-EDGE2 --esg-id edge-ID_GW-EDGE2
```

- c Add the BGP peer to the BGP speaker.

```
neutron bgp-speaker-peer-add name_bgp_speaker name_GW-EDGE1
```

```
neutron bgp-speaker-peer-add name_bgp_speaker name_GW-EDGE2
```

- d Associate the speaker with the VXLAN network.

```
neutron bgp-speaker-network-add name_bgp_speaker external_VXLAN_network_name
```

11 (Optional) Create BGP routers for tenants.

Tenant users can create their BGP routers. The tenant user must be admin to configure a router without SNAT.

- a Create two logical switches for a tenant and subnet pools for them.

```
neutron net-create name_Tenant1_LS1
```

```
neutron subnet-create --name name_network_Tenant1-LS1 name_Tenant1_LS1 --subnetpool selfservice
```

```
neutron net-create name_Tenant1_LS2
```

```
neutron subnet-create --name name_network_Tenant1-LS2 name_Tenant1_LS2 --subnetpool selfservice
```

- b Create a router with BGP configuration.

BGP works with all OpenStack Logical Routers form factors : shared , distributed , and exclusive.

```
neutron router-create name_Tenant1-LR --router_type=exclusive
```

```
neutron router-interface-add name_Tenant1-LR name_network_Tenant1-LS1
```

```
neutron router-interface-add name_Tenant1-LR name_network_Tenant1-LS2
```

```
neutron router-gateway-set name_Tenant1-LR --disable-snat external_VXLAN_network_name
```

BGP dynamic routing is now configured on the provider side and tenants can also use it.

Adding Capacity in the vSphere Web Client

You can add compute clusters and datastores to an existing VMware Integrated OpenStack deployment.

Adding Compute Resources from Multiple vCenter Server Instances

You can add multiple vCenter Server instances to your VMware Integrated OpenStack deployment, if you use VMware NSX-T as virtual network provider. By adding more instances, you improve scalability and resiliency for your OpenStack infrastructure.

Requirements for Adding Compute Resources from Multiple vCenter Server Instances

You can add multiple vCenter Server instances, if the following requirements are met:

- High Availability deployments only, not available for compact mode deployments.
- Virtual network must be provided by VMware NSX-T, not available for deployments that use vSphere Distributed Switch or VMware NSX for vSphere .
- You can have only one availability zone per vCenter Server instance.

Add New Compute Clusters for an OpenStack Deployment from Multiple vCenter Server Instances

You can increase the number of compute clusters in your VMware Integrated OpenStack deployment to increase CPU capacity. You can select compute clusters from all Compute vCenter Server instances in your data center.

Prerequisites

Prepare a cluster with at least one host.

Procedure

- 1 In vCenter Server, select **Home > VMware Integrated OpenStack > Manage**.
- 2 (Optional) Add additional vCenter Server instances for use in VMware Integrated OpenStack.
 - a Select the **Compute vCenter Server** tab.
 - b Click the green plus-sign icon (+) at the top of the panel to add a new instance.
 - c In the Add Compute vCenter Server dialog box, enter the FQDN of the instance, credentials with administrative privileges, and click **OK**.

- 3 Select the **Nova Compute** tab.

This tab displays the current Nova Compute clusters and their status.

- 4 Click the green plus-sign icon (+) at the top of the panel.
- 5 On the Select a Compute vCenter Server page, select the instance and the availability zone for the compute cluster that you need and click **Next**.
- 6 On the Add Nova cluster page, select the cluster that you prepared as a prerequisite, and click **Next**.
The cluster you select must contain at least one host.
- 7 On the Add Nova datastores page, select the datastores for the tenants in the new cluster, and click **Next**.
- 8 On the Review proposed configuration page, select the existing management VM, and click **Next**.
- 9 Review the proposed configuration, and click **Finish**.
- 10 Confirm that the new cluster is added to the OpenStack deployment.
The newly added cluster appears in the **Nova Compute** tab.

OpenStack capacity increases based on the resources available in the additional cluster.

Add Storage to the Compute Node

You can increase the number of datastores available to the Compute node in your VMware Integrated OpenStack deployment.

Adding a datastore to the Compute node causes the Nova service to restart, which might cause a temporary disruption to the OpenStack services in general.

Prerequisites

Verify that you have datastores available. See the vSphere Web Client documentation.

Procedure

- 1 In vCenter, select **Home > VMware Integrated OpenStack > Manage**.
- 2 Click the **Nova Storage** tab.
This tab displays the datastores that are currently available, their status, and other details.
- 3 Click the green plus-sign icon (+) at the top of the panel.
- 4 On the Select a Nova node page of the Add Nova Datastores dialog box, select the cluster to which you want to add a datastore, and click **Next**.
- 5 On the Add Nova datastore page, select one or more datastores to add to the cluster, and click **Next**.
- 6 Review the proposed configuration, and click **Finish**.

The storage capacity for the selected Compute node increases accordingly with the size of the additional datastore.

Add Storage to the Image Service

You can increase the number of datastores available to the Image Service node in your VMware Integrated OpenStack deployment.

Adding a datastore to the Image Service node causes the Glance service to restart, which might cause a temporary disruption to the OpenStack services in general.

Prerequisites

Verify that you have datastores available. See the vSphere Web Client documentation.

Procedure

- 1 In vCenter, select **Home > VMware Integrated OpenStack > Manage**.
- 2 Click the **Glance Storage** tab.
This tab displays the datastores that are currently available, their status, and other details.
- 3 Click the green plus-sign icon (+) at the top of the panel.
- 4 On the Add Glance datastore page, select one or more datastores to add to the cluster, and click **Next**.
- 5 Review the proposed configuration, and click **Finish**.

The storage capacity for the Image Service node increases accordingly with the size of the additional datastore.

Fix Out of Sync Availability Zones

When you work with multiple vCenter Server instances, you might see that availability zones are not synced between the OpenStack management server and the VMware Integrated OpenStack Horizon dashboard.

Due to a known OpenStack issue, if you use CLI commands to rename availability zones, you might see different names in the vSphere Web Client and the Horizon dashboard. In the vSphere Web Client, on the **Nova Compute** tab, under Availability Zones, out of sync availability ones appear in red. Resynchronize the availability zones to fix the issue.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 List the availability zones for your OpenStack deployment.

You need the OpenStack admin password.

```
viocli inventory-admin show-availability-zones
```

3 Synchronize availability zones.

```
viocli inventory-admin sync-availability-zones
```

All availability zones are now synchronized.

Configure the Backup Service for Block Storage

It is a best practice to configure a backup service for the Block Storage (Cinder) component of OpenStack to prevent loss of data. You can configure Cinder to back up volumes to either a network file system (NFS) or an Object Storage (Swift) service, which is another OpenStack service.

You configure a backup service by installing OpenStack Debian packages that are included in your VMware Integrated OpenStack 4.0 deployment.

For the purposes of this procedure, the two controllers are referred to as controller01 and controller02.

Prerequisites

Verify that your VMware Integrated OpenStack 4.0 deployment is installed and running.

For Swift service backup configurations:

- Verify that the Swift component is installed as part of your VMware Integrated OpenStack 4.0 deployment. See the VMware Integrated OpenStack Installation and Configuration Guide.
- Verify that the Swift component is registered to the Identity Service component (Keystone), which is another OpenStack service. This registration is part of the default Keystone configuration. Keystone is installed as part of your VMware Integrated OpenStack 4.0 deployment.

For NFS share backup configurations:

- Create a dedicated NFS share folder to store the backed-up data.
- Verify that the owner of the NFS share folder has the same UID as Cinder on the controller nodes. The default Cinder UID is 107. This value might be different in your deployment.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 3 To use Swift as a backup service, edit the `/opt/vmware/vio/custom/custom.yml` file.
 - a Uncomment the `cinder_backup_driver` parameter.
 - b Set the `cinder_backup_driver` parameter to `cinder.backup.drivers.swift`.

```
# Driver to use for backups. (string value)
cinder_backup_driver: cinder.backup.drivers.swift
```


- 4 To use NFS as a backup service, edit the `/opt/vmware/vio/custom/custom.yml` file.
 - a Uncomment the `cinder_backup_driver` parameter.
 - b Set the `cinder_backup_driver` parameter to `cinder.backup.drivers.nfs`.

```
# Driver to use for backups. (string value)
cinder_backup_driver: cinder.backup.drivers.nfs
```

- c Uncomment the `cinder_backup_share` parameter.
- d Set the `cinder_backup_share` parameter to `<NFS host IP address>:<file backup path>`.

```
# NFS share in fqdn:path, ipv4addr:path, or "[ipv6addr]:path"
# format. (string value)
cinder_backup_share: <NFS host IP address>:<file backup path>
```

- e If your NFS share is not version 4.1, you must uncomment the `cinder_backup_mount_options` parameter and set it to your version of NFS. For example, `vers=3`.

```
# Mount options passed to the NFS client. See NFS man page for
# details. (string value) 'vers=4' to support version NFS 4
cinder_backup_mount_options: vers=4
```

- 5 Save the `custom.yml` file.
- 6 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment -v configure --limit controller
```

Important This command updates your entire deployment and might briefly interrupt operations.

What to do next

Verify that the Cinder backup configuration functions properly. See [Verify That the Cinder Backup Service is Running and Operational](#)

Verify That the Cinder Backup Service is Running and Operational

Create and back up a test volume to verify that the Cinder backup is properly configured and running.

Prerequisites

Complete the Cinder backup configuration. See [Configure the Backup Service for Block Storage](#).

Procedure

- 1 Confirm that the Cinder backup service is running.

```
cinder service-list
```

- 2 Create a test volume.

```
cinder create --display-name <volume name>
```

- 3 Create a backup of the test volume.

```
cinder backup-create --display-name <backup name> <volume name>
```

- 4 Check the NFS share or Swift service to confirm that the backup file was created.

Troubleshoot Cinder Volume Backup Failure

While you are configuring the Cinder backup on an NFS share, the first attempt to create a test backup fails.

Problem

When you verify the Cinder backup configuration, you get an error when you create the initial backup.

Cause

VMware Integrated OpenStack does not have the correct permissions to write to the NFS share.

Solution

- 1 Using SSH, log in to the controller01 node as the root user.
- 2 Go to the mount directory for the Cinder backup configuration.

```
cd /var/lib/cinder/backup_mount/
```

- 3 Change the folder owner from root to cinder.

```
chown -R cinder:cinder *
```

This workaround corrects the configuration and gives the Cinder component permission to access the NFS share.

Back Up the VMware Integrated OpenStack Deployment

It is a best practice to periodically back up your OpenStack management server and database.

You perform backup operations in the CLI for the VMware Integrated OpenStack Manager.

Prerequisites

You must log in with administrative or super-user (sudo) privileges to perform backup operations.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.

- 2 Switch to root user.

```
sudo su -
```

- 3 (Optional) Switch to verbose mode.

```
viocli backup <-v | -verbose>
```

- 4 (Optional) View the help options.

```
viocli backup <-h | -help>
```

- 5 Use the `viocli backup mgmt_server <NFS_VOLUME>` command to back up the OpenStack management server.

```
viocli backup mgmt_server [-d DEPLOYMENT] <NFS_VOLUME>
```

Option	Description
<code>-d DEPLOYMENT</code>	Specifies the name of the VMware Integrated OpenStack deployment to be backed up.
<code>NFS_VOLUME</code>	Name or IP address of the target NFS volume and directory in the format <i>remote_host:remote_dir</i> . For example: 192.168.1.77:/backups

The backup file is automatically labeled with the timestamp `vio_ms_yyyymmddhhmmss`.

- 6 Back up the OpenStack database.

```
viocli backup openstack_db [-d DEPLOYMENT] <NFS_VOLUME>
```

Option	Description
<code>-d DEPLOYMENT</code>	Specifies the name of the VMware Integrated OpenStack deployment database to be backed up..
<code>NFS_VOLUME</code>	Name or IP address of the target NFS volume and directory in the format <i>remote_host:remote_dir</i> . For example: 192.168.1.77:/backups

The backup file is automatically labeled with the timestamp `vio_os_db_yyyymmddhhmmss`.

If a severe event occurs, you can use the new backup files to restore your VMware Integrated OpenStack deployment data and configuration.

Restore VMware Integrated OpenStack from a Backup

If a crash occurs, you can restore your VMware Integrated OpenStack management server and OpenStack database from a previous backup.

You perform restore operations in the CLI for the VMware Integrated OpenStack Manager.

Prerequisites

Log in with administrative or super-user (sudo) privileges to perform restore operations.

Verify that you have backups of the management server and database available. See [Back Up the VMware Integrated OpenStack Deployment](#).

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 Switch to root user.

```
sudo su -
```

- 3 (Optional) Switch to verbose mode.

```
viocli restore <-v | -verbose>
```

- 4 (Optional) View the help options.

```
viocli restore <-h | -help>
```

- 5 Restore the OpenStack management server, where PATH specifies the intended location for the backup file..

```
viocli restore mgmt_server \  
[-d DEPLOYMENT] \  
<BACKUP_NAME> \  
<NFS_VOLUME>
```

Option	Description
-d DEPLOYMENT	Indicates the backup by the deployment name assigned when it was created.
BACKUP_NAME	Indicates the timestamp label of the backup file to be used to restore the management server.
NFS_VOLUME	Indicates the NFS host where the backup file is located.

- 6 Restore the OpenStack database.

```
viocli restore openstack_db \  
[-d DEPLOYMENT] \  
<BACKUP_NAME> \  
<NFS_VOLUME>
```

Option	Description
-d DEPLOYMENT	Indicates the backup by the deployment name assigned when it was created.
BACKUP_NAME	Indicates the timestamp label of the backup file to be used to restore the database.
NFS_VOLUME	Indicates the NFS host where the backup file is located.

You restore your VMware Integrated OpenStack management server and OpenStack database to the state of the backups.

Failure Recovery

In the event of a disk failure or another critical issue, you can recover the individual nodes in your VMware Integrated OpenStack deployment using the CLI.

When you recover a VMware Integrated OpenStack node, it returns to the state of a newly deployed node. To recover a database node, you must recover to a backup file. See [Back Up the VMware Integrated OpenStack Deployment](#).

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 Switch to root user.

```
sudo su -
```

- 3 Switch to verbose mode.

```
viocli recover <-v | -verbose>
```

- 4 View the help options.

```
viocli recover <-h | -help>
```

5 Recover the OpenStack nodes by node or role.

a To recover a database node:

```
viocli recover <[-r ROLE -dn BACKUP_NAME] | [-n NODE -dn BACKUP_NAME]> -nfs NFS_VOLUME
```

Option	Description
-n NODE	<p>Recovers the database nodes specified by VM name recover by node name. You can specify multiple nodes in one command.</p> <p>Use the VM name as it appears in the VMware Integrated OpenStack manager (VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name]).</p> <p>For example,</p> <pre>viocli recover -n VIO-DB-0 VIO-DB-1 VIO-DB-2 -dn vio_os_db_20150830215406 -nfs 10.146.29.123:/backups</pre> <p>recovers from the specified NFS backup file all the named database nodes: VIO-DB-0, VIO-DB-1, and VIO-DB-2.</p>
-r ROLE	<p>Recovers all the database nodes in the specified group name. You can specify multiple roles in one command.</p> <p>Use the group name as it appears in the VMware Integrated OpenStack manager (VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name]).</p> <p>For example,</p> <pre>viocli recover -r DB -dn vio_os_db_20150830215406 -nfs 10.146.29.123:/backups</pre> <p>recovers from the specified NFS backup file all the nodes in the DB node group.</p>
-dn BACKUP_NAME	Indicates the timestamp label of the backup file to be used to restore the database.
-nfs NFS_VOLUME	Indicates the NFS host where the backup file is located.

b To recover any non-database node:

```
viocli recover <[-r ROLE] | [-n NODE]>
```

Option	Description
-n NODE	<p>Recovers the nodes specified by VM name. You can specify multiple nodes in one command.</p> <p>Use the VM name as it appears in the VMware Integrated OpenStack manager (VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name]).</p> <p>For example,</p> <pre>viocli recover -n VIO-Controller01</pre>

Option	Description
	recovers the VIO-Controller01 node.
<code>-r ROLE</code>	<p>Recovers all the nodes in the specified group name. You can specify multiple roles in one command.</p> <p>Use the group name as it appears in the VMware Integrated OpenStack manager (VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name]).</p> <p>For example,</p> <pre>viocli recover -r VIO-Controller01</pre> <p>recovers all nodes in the VIO-Controller01 node group.</p>



Tip You can use the `viocli show` command to list all the nodes and their roles in your VMware Integrated OpenStack deployment.

- Verify the node is running by checking its status in the VMware Integrated OpenStack manager: **VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name]**.

Depending on your deployment, the recovery process might take a few minutes.

VMware Integrated OpenStack Log File Locations

When you request technical support, you might be requested to provide log files. The following tables show you where the files are located and describes their purpose.

VMware Integrated OpenStack Management Server Logs

Name and Location	Description
<code>/var/log/apache2/access.log</code>	Logs access to the VMware Integrated OpenStack Manager.
<code>/var/log/apache2/error.log</code>	Logs access errors for the VMware Integrated OpenStack Manager.
<code>/var/log/jarvis/ansible.log</code>	Logs Ansible service activity.
<code>/var/log/jarvis/jarvis.log</code>	Logs Jarvis service activity.
<code>/var/log/jarvis/pecan.log</code>	Logs Pecan framework service activity.
<code>/var/log/oms/oms.log</code>	Logs VMware Integrated OpenStack Manager service activity.
<code>/var/log/oms/register-plugin.log</code>	Logs VMware Integrated OpenStack plugin registration activity.
<code>/var/log/osvmw/osvmw-exceptions.log</code>	Logs exceptions to osvmw service.
<code>/var/log/osvmw/osvmw.log</code>	Logs osvmw service activity.
<code>/var/log/viocli/viocli.log</code>	Logs viocli (VMware Integrated OpenStack CLI) service activity.
<code>/var/log/viomon/viomon.log</code>	Logs VMware Integrated OpenStack monitoring activity.
<code>/var/log/viopatch/*.log</code>	Logs upgrade and patching activity.
<code>/var/log/bootsequence.log</code>	Logs booting activity.

OpenStack Controller Logs

Name and Location	Description
/var/log/apache2/access.log	Logs Horizon (VMware Integrated OpenStack dashboard) access activity.
/var/log/cinder/cinder-api.log	Logs Cinder API service activity.
/var/log/apache2/error.log	Logs Horizon (VMware Integrated OpenStack dashboard) general activity.
/var/log/cinder/cinder-scheduler.log	Logs Cinder Scheduler service activity.
/var/log/glance/glance-api.log	Logs Glance API service activity.
/var/log/cinder/cinder-volume.log	Logs Cinder volume service activity.
/var/log/glance/glance-registry.log	Logs Glance registry service activity.
/var/log/glance/manage.log	Logs Glance service general activity.
/var/log/heat/heat-api-cfn.log	Logs Heat service general activity.
/var/log/heat/heat-api-cloudwatch.log	Logs Heat service general activity.
/var/log/heat/heat-api.log	Logs Heat API service activity.
/var/log/heat/heat-engine.log	Logs Heat engine service activity.
/var/log/keystone/keystone-manage.log	Logs Keystone manage service activity.
/var/log/keystone/keystone.log	Logs Keystone service general activity.
/var/log/neutron/neutron-server.log	Logs Neutron server service activity.
/var/log/nova/nova-api.log	Logs Nova API service activity.
/var/log/nova/nova-conductor.log	Logs Nova conductor service activity.
/var/log/nova/nova-consoleauth.log	Logs Nova consoleauth service activity.
/var/log/nova/nova-manage.log	Logs Nova manage service activity.
/var/log/nova/nova-mksproxy.log	Logs Nova mksproxy service activity.
/var/log/nova/nova-novncproxy.log	Logs Nova novncproxy service activity.
/var/log/nova/nova-scheduler.log	Logs Nova scheduler service activity.

Database Service Logs

Name and Location	Description
/var/log/syslog	General database logging including MySQL logging.
/var/log/rabbitmq/rabbit@database01.log	Logs general RabbitMQ database activity.
/var/log/rabbitmq/shutdown_log	Logs RabbitMQ service shut-down activity.
/var/log/rabbitmq/startup_log	Logs RabbitMQ service start-up activity.

Compute and Loadbalancer Service Logs

Name and Location	Description
/var/log/haproxy/haproxy.log	Logs HAProxy service activity.
/var/log/nova/nova-compute.log	Logs Nova compute service activity.
/var/log/nova/nova-manage.log	Logs Nova manager service activity.
/var/log/nova/vmware-vspc.log	Logs VMware Virtual Serial Port Concentrator (VSPC) activity.
/var/log/ceilometer/ceilometer-agent-compute.log	Logs Ceilometer agent activity.

Upgrade to VMware Integrated OpenStack 4.0

You upgrade to VMware Integrated OpenStack 4.0 by deploying a separate instance, setting up an NFS server where you back up the current deployment, reconfiguring, and migrating to the new, upgraded deployment.

Important You can only upgrade to VMware Integrated OpenStack 4.0 from VMware Integrated OpenStack 3.1. If you are running a version different from 3.1, you must first upgrade to that version. For more information about upgrading to 3.1, see the [Upgrade to VMware Integrated OpenStack 3.0 or 3.1](#) procedure.

The update process requires vSphere to accommodate the existing deployment and the upgraded deployment. You must make available additional resources, datastores, IP addresses, to complete the upgrade procedure. vSphere continues to hold both deployments until you determine that the upgrade process was successful and you do not need to roll back to your previous VMware Integrated OpenStack deployment.

Important Upgrade only preserves customizations configured in the `custom.yml` file. Any changes or customizations made directly to the OpenStack deployment, such as SWIFT, are not preserved. It is responsibility of the OpenStack administrator to track such changes and re-apply them after the upgrade.

Prerequisites

- Download the latest VMware Integrated OpenStack 4.0 OVA from the VMware website.
- Verify that you have matching resources for every node except the memcache and RabbitMQ nodes. See the hardware requirements in the VMware Integrated OpenStack Installation and Configuration Guide.

Procedure

1 [Add IP Addresses to the Network Configuration](#)

The upgrade procedure requires a temporary number of IP addresses in addition to your existing IP address configuration. vSphere provides a tool that enables you to add this required IP range.

2 [Deploy a VMware Integrated OpenStack 4.0 Instance, Backup the Existing, and Migrate Your Management Server Data](#)

To upgrade to VMware Integrated OpenStack 4.0, first deploy a new instance with the latest OVA. You use the new deployment to set up a NFS server and backup your previous management server instance on the new one.

3 [Migrate to the VMware Integrated OpenStack 4.0 Deployment](#)

After you prepare the NFS server and backup your current deployment, you install it as a separate deployment and migrate your data.

4 [Revert to a Previous VMware Integrated OpenStack Deployment](#)

You can revert to VMware Integrated OpenStack to a previous version by restoring your previous deployment.

5 [Delete the Older VMware Integrated OpenStack Deployment](#)

After you complete the upgrade process to the VMware Integrated OpenStack 4.0 deployment, you can delete the older VMware Integrated OpenStack deployment. By deleting the old deployment, you recover the CPU, datastores, and IP addresses resources that it required.

Add IP Addresses to the Network Configuration

The upgrade procedure requires a temporary number of IP addresses in addition to your existing IP address configuration. vSphere provides a tool that enables you to add this required IP range.

You can use this procedure to add IP addresses for any reason. If you are not adding IP addresses as part of the upgrade procedure, the specific number of IP addresses required might not apply.

Procedure

1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.

2 Click the **Manage** tab and click the **Networks** tab.

The **Networks** tab lists the Management and API network configurations, including their IP address ranges.

3 Expand the IP addresses available for the Management network.

a Right-click the name of the Management network in the list and select **Add IP Range**.

b In the Add IP Range dialog box, specify the new IP range.

Note If you are adding addresses as part of the upgrade process, the new IP range must match the same number of IP addresses configured for the existing Management network. For example, in a typical VMware Integrated OpenStack deployment, the Management network requires a minimum range of 11 IP addresses.

c Click **OK**.

- 4 Expand the IP addresses available for the external network.
 - a Right-click the name of the API network in the list and select **Add IP Range**.
 - b In the Add IP Range dialog box, specify the new IP range.

Note If you are adding addresses as part of the upgrade process, the new IP range must match the same number of IP addresses configured for the existing API network. For example, in a typical VMware Integrated OpenStack deployment, the API network requires a minimum range of 2 IP addresses for versions 3.0 and older and a minimum range of 3 IP addresses for versions 3.1 and later.

- c Click **OK**.

What to do next

If you added IP addresses as part of the upgrade process, you can now obtain and install the upgrade patch.

Deploy a VMware Integrated OpenStack 4.0 Instance, Backup the Existing, and Migrate Your Management Server Data

To upgrade to VMware Integrated OpenStack 4.0, first deploy a new instance with the latest OVA. You use the new deployment to set up a NFS server and backup your previous management server instance on the new one.

Procedure

- 1 Deploy a VMware Integrated OpenStack 4.0 instance.

For detailed instructions, see *Installing Integrated OpenStack* in the *VMware Integrated OpenStack Installation and Configuration Guide*.

- 2 On the new VMware Integrated OpenStack management server, set up an NFS server and copy the SSH key of the previous instance.
 - a Log in to the new VMware Integrated OpenStack management server over SSH.
 - b Run the following command to create a directory for use by the NFS server to store the backup of the previous VMware Integrated OpenStack management server.

```
sudo viocli upgrade prepare 3.1_oms_ip /folder_for_nfs_server
```

For example: `sudo viocli upgrade prepare 192.168.100.101 /data`

- c Enter the password for the *viouser* account and wait for the operation to finish.
 - d After the operation finishes, keep the console window open.

- 3 Backup the data from your VMware Integrated OpenStack management server by using the configured NFS server on the new instance.
 - a Open another session to log in to the previous VMware Integrated OpenStack 3.1 management server over SSH.
 - b Backup the data of the previous VMware Integrated OpenStack management server on the new one.

You can also use the optional verbose parameter.

```
sudo viocli backup mgmt_server 4.0_oms_ip:/folder_for_nfs_server
```

For example: `sudo viocli backup mgmt_server 192.168.100.102:/data --verbose`

- 4 Migrate the backed up data from your previous VMware Integrated OpenStack management server to the new one.
 - a Return to the console connected to the new VMware Integrated OpenStack management server.
 - b List the contents of the folder created for use by the NFS server.

For example:

```
cd /data
ll
```

- c Copy the name of the directory that begins with `vio_ms_`.
- d Reconfigure the new instance by using the folder name from the previous step, the IP address of the new server, and the NSF server folder.

You can also use the optional verbose parameter.

```
sudo viocli upgrade mgmt_server folder_containing_backup 4.0_oms_ip:/folder_for_nfs_server
```

For example: `sudo viocli upgrade mgmt_server vio_ms_20170918093000 192.168.100.102:/data --verbose`

The VMware Integrated OpenStack management server is now upgraded.

What to do next

You can now install and provision the new VMware Integrated OpenStack vApp.

Migrate to the VMware Integrated OpenStack 4.0 Deployment

After you prepare the NFS server and backup your current deployment, you install it as a separate deployment and migrate your data.

The upgrade process also allows you to switch from a compact mode deployment to HA deployment.

Procedure

- 1 If you are logged in to the vSphere Web Client, log out and log back in.

This refreshes the interface so the newly installed deployment is accessible through the vSphere Web Client.

- 2 In the vSphere Web Client, click **Home** and click the VMware Integrated OpenStack icon.
- 3 Click the **Summary** tab and verify that the **Version information** table shows the version of the upgraded VMware Integrated OpenStack manager.
- 4 Click the **Manage** tab and click the **Upgrades** tab.

The **Upgrades** tab lists the current VMware Integrated OpenStack deployment.

- 5 Right-click the deployment name and select **Upgrade**.
- 6 Enter name for the new deployment.

This name must be different than the name of the existing deployment.

- 7 If you are upgrading from a compact mode deployment, in the **Deployment type** drop-down menu, select type for the upgraded deployment.

If your VMware Integrated OpenStack is deployed in compact mode, you can change to HA mode during the update or stay in compact mode.

- 8 Click **Next**.
- 9 Review the upgrade configuration, and click **Finish**.

The current deployment shows a status of **Running** and the new, upgraded deployment shows a status of **Prepared**.

- 10 On the **Upgrades** tab, right-click the name of the old deployment, and select **Migrate Data**.

Important You must confirm this action because during data migration, the VMware Integrated OpenStack services stop and downtime incurs until the upgrade finishes.

When the migration process finishes, the status for the updated deployment on the **Upgrades** tab changes to **Migrated**. The new deployment is up and running, but it is using a temporary public IP. Users can access it, but only using this temporary public VIP. Only after completion of the following step, the original public VIP is configured on the new deployment. The **Upgrades** tab now lists the current VMware Integrated OpenStack and the new deployment.

- 11 On the **Upgrades** tab, right-click the name of your previous deployment, and select **Switch to New Deployment**.

When the deployment switching process finishes, the status for the updated deployment on the **Upgrades** tab changes to **Running**. The previous deployment shows a status of **Stopped**.

What to do next

If the deployment process is unsuccessful, you can revert to your previous VMware Integrated OpenStack deployment. See [Revert to a Previous VMware Integrated OpenStack Deployment](#).

If the deployment process is successful, you can delete the previous VMware Integrated OpenStack deployment.

Revert to a Previous VMware Integrated OpenStack Deployment

You can revert to VMware Integrated OpenStack to a previous version by restoring your previous deployment.

Prerequisites

- Verify that you retained the previous VMware Integrated OpenStack deployment in your OpenStack manager.
- Verify that you are prepared to stop the services running on the previous VMware Integrated OpenStack deployment.

Procedure

- 1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.
- 2 In the Inventory Lists panel, click **OpenStack Deployments**.
The current VMware Integrated OpenStack deployment appears in the center pane.
- 3 Right-click the name of the current deployment on the **OpenStack Deployments** tab and select **Stop OpenStack Deployment**.
- 4 Return to the main VMware Integrated OpenStack panel (**Home > Inventories > VMware Integrated OpenStack**).
- 5 Click the **Manage** tab and click the **Upgrades** tab.
The **Upgrades** tab lists the VMware Integrated OpenStack 4.0 and older deployments.
- 6 Right-click the previous VMware Integrated OpenStack deployment name and select **Restore** from the pop-up menu.

When the process of reverting your VMware Integrated OpenStack deployment is finished, the OpenStack services restart.

Delete the Older VMware Integrated OpenStack Deployment

After you complete the upgrade process to the VMware Integrated OpenStack 4.0 deployment, you can delete the older VMware Integrated OpenStack deployment. By deleting the old deployment, you recover the CPU, datastores, and IP addresses resources that it required.

Prerequisites

Verify that the your upgraded VMware Integrated OpenStack 4.0 deployment is running and functioning successfully. After you delete a deployment, you cannot restore it.

Procedure

- 1 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.
- 2 Click the **Manage** tab and click the **Upgrades** tab.

The **Upgrades** tab lists the current and old VMware Integrated OpenStack and old deployments. The VMware Integrated OpenStack 4.0 deployment shows a status of Running. The previous VMware Integrated OpenStack deployment shows a status of Stopped.

- 3 Right-click the older VMware Integrated OpenStack deployment and select **Delete** from the pop-up menu.
- 4 At the prompt, confirm the deletion.

The deployment no longer appears on the **Upgrades** tab or in the **OpenStack Deployments** list.

Updating Your VMware Integrated OpenStack Deployment

You update your VMware Integrated OpenStack deployment by using the VMware Integrated OpenStack Manager vApp or CLI commands to install and apply patches.

After installing a patch, you can revert to a previous version if necessary.

Install Patch Using the vSphere Web Client

VMware provides updates in the form of Debian patches. Patches that do not affect the infrastructure of the VMware Integrated OpenStack deployment can be applied by using the VMware Integrated OpenStack Manager vApp.

Prerequisites

vSphere Web Client

Some patches might require you to shut down the VMware Integrated OpenStack service before proceeding.

Procedure

- 1 Download the Debian patch from VMware.

If you do not know where to obtain the patch, go to the VMware Integrated OpenStack product page <https://www.vmware.com/products/openstack> or consult with VMware.

- 2 Transfer the patch file to the management server.

- 3 Log in to the management server and enter the following command to load the patch file into the management server repository:

```
sudo viopatch add -l path/filename.deb
```

where *filename.deb* is the filename of the Debian patch file.

- 4 In the vSphere Web Client, select **Home > Inventories**, and click the VMware Integrated OpenStack icon.

- 5 Click the **Manage** tab and click the **Updates** tab.

The **Updates** tab lists added patches and indicates if they are installed.

- 6 Select the patch and click **Choose**.

The patch appears in the list on the **Updates** tab.

- 7 Install the patch.

If you can install the patch by using the VMware Integrated OpenStack Manager vApp, the **Apply** option appears in the Patch Action column on the **Updates** tab.

If the **Apply** option does not appear in the Patch Action column, click **More details** in the Patch Description column to access instructions for installing patches by using the CLI.

After you install a patch, the value in the Patch Status column on the **Updates** tab changes to Installed.

- 8 To complete the update, log out of the vSphere Web Client and back in.

You can ignore any error messages you encounter when logging back in.

- 9 Restart all VMware Integrated OpenStack services.

Install Patch Using CLI Commands

VMware provides updates in the form of Debian patches. Patches that affect the infrastructure of the VMware Integrated OpenStack deployment must be applied through the command console for the VMware Integrated OpenStack Manager vApp.

Procedure

- 1 Download the Debian patch from VMware.

If you do not know where to obtain the patch, go to the VMware Integrated OpenStack product page <https://www.vmware.com/products/openstack> or consult with VMware.

- 2 Add the patch to your VMware Integrated OpenStack installation.
 - a Log into the console for the VMware Integrated OpenStack management server.
 - b Add the patch.

```
viopatch add -l [path to the debian file]
```

- c Confirm that the patch was successfully added.

```
viopatch list
```

This returns a list of available patches, their version numbers, their type, and current status. The list should indicate the patch by its build number.

- 3 Install the patch.
 - a Ensure that VMware Integrated OpenStack service is either running or not yet deployed. If the VMware Integrated OpenStack service is in any other state, the upgrade will fail.
 - b Log into the VMware Integrated OpenStack management server and run the following command:

```
viopatch install -p <upgrade patch name> -v <upgrade patch version>
```

The patch installation takes 5 to 10 minutes to complete.

- 4 To complete the update, log out of the vSphere Web Client and back in. You can ignore any error messages you encounter when logging back in.
- 5 Restart all VMware Integrated OpenStack services.

If necessary, you can revert to the previous version. For details, see [Reverting a Patch Update Installation](#).

For troubleshooting the patch installation, see [Troubleshooting Update Patch Installation Problems](#)

Reverting a Patch Update Installation

You can revert a patch update installation.

Prerequisites

You can only revert to an earlier version of the same point release. For example, you cannot revert a 2.0 implementation to a 1.0.x version.

Procedure

- 1 Log into the console for the VMware Integrated OpenStack management server.

- 2 Run the uninstall command.

```
viopatch uninstall --patch vio-patch-[version number] --version [build number]
```

The reversion process takes 5 to 10 minutes to complete.

- 3 After uninstalling the patch, restart the vSphere Web Client service on the vCenter Server to downgrade the VMware Integrated OpenStack plugin.

Troubleshooting Update Patch Installation Problems

This section describe some common errors you might encounter while installing the update patch.

Troubleshoot Update Patch Installation Failure

The patch installation fails.

Problem

After adding and applying the update patch, the installation fails.

Cause

The VMware Integrated OpenStack deployment must be running or not yet deployed.

Solution

- 1 Ensure that the VMware Integrated OpenStack service is either running or not yet deployed.
- 2 If the service is running, ensure that all the OpenStack management VMs (database, load balancer, and so on) are also running.

Troubleshoot Update Patch Installation Errors

You get an error when using the vSphere Web Client to add patch.

Problem

Patch installation fails with a fatal error message in vSphere Web Client.

Cause

The type of update requires using the CLI to add and install the patch

Solution

- ◆ Add and install the patch using the CLI method described in [Install Patch Using CLI Commands](#).

Customize the Dashboard Logos and Background

By default, the VMware Integrated OpenStack dashboard log-in page displays the VMware corporate logo and a blank background. All pages in the VMware Integrated OpenStack dashboard display the VMware corporate logo in the upper left hand corner. You can customize your deployment configuration to display your company's logos or other branding instead of these default graphics.

- The default dimension for the logo graphic is 216 pixels long by 35 pixels wide. You can use a graphic with different dimensions, but the display might be impacted.
- The background graphic appears in the center of the log-in page.

Procedure

1 [Customize the Log-in Page Background](#)

You can specify a custom graphic to appear as the background to the VMware Integrated OpenStack dashboard log-in page.

2 [Customize the Log-in Page Logo](#)

You can specify the custom logo that appears on the VMware Integrated OpenStack dashboard log-in page.

3 [Customize the Dashboard Page Logo](#)

You can specify the custom logo that appears in the top left corner of each page in the VMware Integrated OpenStack dashboard.

Customize the Log-in Page Background

You can specify a custom graphic to appear as the background to the VMware Integrated OpenStack dashboard log-in page.

Procedure

- 1 Load your custom graphic file in the `/home/viouser/custom/horizon/` directory in your VMware Integrated OpenStack deployment.

This directory is the default directory for graphic files in the VMware Integrated OpenStack dashboard.

2 Open the `/home/viouser/custom/horizon/_styles.scss` file in a text editor.

a Uncomment the `.login-bg` parameter.

```
.login-bg {
  height: 100%;
  body {
    background: #1D2226 url("/static/themes/vmware/CUSTOM-BACKGROUND-IMAGE.jpg") repeat-x 45%
0 !important;
    background-size: 100% auto !important;
    color: black;
  }
}
```

b Modify the `.login-bg` parameter to reference your custom background graphic file.

c Save the `_styles.scss` file.

3 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

4 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.

a Uncomment the parameter that enables the `custom.yml` settings to override the default style sheet settings.

```
# overwrite the _styles.scss file in the VMware theme
horizon_custom_stylesheet: "/home/viouser/custom/horizon/_styles.scss"
```

b Uncomment the parameter that specifies the custom directory to contain the custom graphic file.

```
# copy all custom images (or other files) to be accessible in horizon
# IMPORTANT: this line must end with a "/" in order to place the files
# in the right location for horizon
horizon_custom_directory: "/home/viouser/custom/horizon/"
```

c Save the `custom.yml` file.

Your custom background image appears on the dashboard log-in page the next time you start a session.

Customize the Log-in Page Logo

You can specify the custom logo that appears on the VMware Integrated OpenStack dashboard log-in page.

Procedure

- 1 Load your custom graphic file to the `/home/viouser/custom/horizon/` directory in your VMware Integrated OpenStack deployment.

This directory is the default directory for graphic files in the VMware Integrated OpenStack dashboard.

- 2 Modify the `/home/viouser/custom/horizon/_styles.scss` file in a text editor.
 - a Uncomment the `.login` parameter.

```
.login {
    background-image: url(/static/themes/vmware/CUSTOM_LOGIN_PAGE_LOGO.png);
    color: white;
    background-color: black;
}
```

- b Modify the `.login` parameter to reference your custom graphic file.
 - c Save the `_styles.scss` file.
- 3 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 4 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
 - a Uncomment the parameter that enables the `custom.yml` settings to override the default style sheet settings.

```
# overwrite the _styles.scss file in the VMware theme
horizon_custom_stylesheet: "/home/viouser/custom/horizon/_styles.scss"
```

- b Uncomment the parameter that specifies the custom directory to contain the custom graphic file.

```
# copy all custom images (or other files) to be accessible in horizon
# IMPORTANT: this line must end with a "/" in order to place the files
# in the right location for horizon
horizon_custom_directory: "/home/viouser/custom/horizon/"
```

- c Save the `custom.yml` file.

Your custom logo appears on the dashboard log-in page the next time you start a session.

Customize the Dashboard Page Logo

You can specify the custom logo that appears in the top left corner of each page in the VMware Integrated OpenStack dashboard.

Procedure

- 1 Load your custom graphic file to the `/home/viouser/custom/horizon/` directory in your VMware Integrated OpenStack deployment.

This directory is the default directory for graphic files in the VMware Integrated OpenStack dashboard.

- 2 Modify the `/home/viouser/custom/horizon/_styles.scss` file in a text editor.
 - a Uncomment the `.topbar` parameter.

```
.topbar {
  h1.brand a {
    background-image: url(/static/themes/vmware/CUSTOM_PAGE_LOGO.png);
  }
}
```

- b Modify the `.topbar` parameter to reference your custom graphic file.
 - c Save the `_styles.scss` file.
- 3 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 4 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
 - a Uncomment the parameter that enables the `custom.yml` settings to override the default style sheet settings.

```
# overwrite the _styles.scss file in the VMware theme
horizon_custom_stylesheet: "/home/viouser/custom/horizon/_styles.scss"
```

- b Uncomment the parameter that specifies the custom directory to contain the custom graphic file.

```
# copy all custom images (or other files) to be accessible in horizon
# IMPORTANT: this line must end with a "/" in order to place the files
# in the right location for horizon
horizon_custom_directory: "/home/viouser/custom/horizon/"
```

- c Save the `custom.yml` file.

Your custom logo appears in the top left corner of each dashboard page the next time you start a session.

Use Profiling to Trace OpenStack Deployments

By using the VMware Integrated OpenStack profiling feature, you can enable tracing for the core OpenStack services. When enabled, tracing captures the response time of all API, RPC, driver, and database calls that are part of an OpenStack operation. You can enable or disable tracing without having to restart OpenStack services.

VMware Integrated OpenStack provides two options for configuring profiler. You can use it either with the Ceilometer OpenStack service or with vRealize Log Insight to store profiler trace data.

Procedure

1 [Configure Tracing of OpenStack Services](#)

Configure the VMware Integrated OpenStack profiling feature by modifying the `custom.yaml` file.

2 [Use Tracing of OpenStack Services](#)

Use the VMware Integrated OpenStack profiling to capture the response time of all API, RPC, driver, and database calls that are part of an OpenStack operation.

Configure Tracing of OpenStack Services

Configure the VMware Integrated OpenStack profiling feature by modifying the `custom.yaml` file.

VMware Integrated OpenStack provides two options for configuring profiler. You can use it either with the Ceilometer OpenStack service or with vRealize Log Insight to store profiler trace data.

Prerequisites

- To use vRealize Log Insight to store profiler trace data, verify that your instance is fully operational, version 3.3 or later, and that you can authenticate with a user with the USER role assigned.
- To use Ceilometer OpenStack service to store profiler trace data, verify that the service is running.

Procedure

1 Modify the custom.yml file to enable tracing.

- a If you have not already done so, implement the custom.yml file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- b Edit the custom.yml file by uncommenting and modifying parameters.

- ◆ If you use Ceilometer OpenStack uncomment and modify the following parameters.

```
os_profiler_enabled: True
os_profiler_hmac_keys: SECRET_KEY
```

- ◆ If you use vRealize Log Insight, uncomment and modify the following parameters.

```
os_profiler_enabled: True
os_profiler_hmac_keys: SECRET_KEY
os_profiler_connection_string:
"loginsight://loginsight_username:password@loginsight_ip_address"
```

Parameter	Description
os_profiler_enabled	Accept the default value. When set to True , the OpenStack profiling feature is enabled.
os_profiler_hmac_keys	Specify the security key. This key must be provided each time an administrator runs a trace.
os_profiler_connection_string	Specify the authentication for the vRealize Log Insight server. Include user name, password and address of the instance.

2 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment configure
```

Note Pushing the configuration briefly interrupts OpenStack services.

- 3 If you use vRealize Log Insight to store profiler trace data, set environment variable `OSPROFILER_CONNECTION_STRING` so that you don't enter connection string each time you run commands with profiling enabled.

You must set the variable on all VMware Integrated OpenStack controllers that you want to run commands from.

```
export
OSPROFILER_CONNECTION_STRING="loginsight://loginsight_username:password@loginsight_ip_address"
```

You can now use the profiling feature.

Use Tracing of OpenStack Services

Use the VMware Integrated OpenStack profiling to capture the response time of all API, RPC, driver, and database calls that are part of an OpenStack operation.

VMware Integrated OpenStack currently supports profiling of Cinder, Heat, Glance, Nova, and Neutron commands.

Prerequisites

- Make sure that you've set environment variable `OSPROFILER_CONNECTION_STRING` on the controller where you will trace the OpenStack services. See, [Configure Tracing of OpenStack Services](#)

Procedure

- 1 Enable profiling by specifying the `profile` option for a given command and provide the secret key.

```
cinder --profile YOUR_SECRET_KEY list
```

The output shows a command that you use to generate the profiling report in HTML format.

- 2 Run the generated command from the output to generate a report, for example `trace.html`.

```
osprofiler trace show --html <UUID> > trace.html
```

For more information on the different options for the report, see the `osprofiler trace show` command help.

```
osprofiler trace show --help
```

Configure NUMA for Use With VMware Integrated OpenStack

VMware Integrated OpenStack 4.0 supports NUMA aware placement on the underlying vSphere platform. This feature provides low latency and high throughput to Virtual Network Functions (VNFs) that run on Telco environments.

For more information about NUMA, see [Using NUMA Instances with ESXi](#).

To achieve low latency and high throughput, it is important that vCPUs, memory, and physical NICs that are used for VM traffic are aligned on same NUMA node. You must create a specific teaming policy that depends on the type of deployment that you have.

In VMware Integrated OpenStack 4.0, selecting physical NIC for placement is a manual configuration.

Procedure

- 1 Login to the ESXi hosts in your data center and run the following command to gather information about the current NUMA configuration.

```
vsish -e get /net/pNics/vmnic<id>/properties | grep 'Device NUMA Node'
```

- 2 If you use overlay networks, all VTEPs are associated to a dvportgroup and you must group all physical NICs on a given NUMA node and create a teaming policy that includes only the physical NICs on that NUMA node.

- 3 If you use VLAN based network backed by dvportgroup, create teaming policy that uses only physical NICs from the given NUMA node and create a Neutron portgroup type provider network.

```
neutron net-create --provider:network_type portgroup <numa_network_name>
```

- 4 Specify the `numa.nodeAffinity` metadata in a VMware Integrated OpenStack flavor.

```
nova flavor-key <uuid> set vmware:extra_config='{"numa.nodeAffinity": "node_ID"}
```

- 5 Boot the instance on OpenStack with this flavor and select a Neutron network where the physical NICs are from same NUMA node.

Managing OpenStack Projects and Users

3

In VMware Integrated OpenStack, cloud administrators manage permissions through user, group, and project definitions. Projects in OpenStack equate to tenants in vCloud Suite. You can assign users and user groups to more than one project.

Before you can create a user, you must create at least one project to which you can assign the user.

This section includes the following topics:

- [Create an OpenStack Project](#)
- [Modify a Project](#)
- [Working with Security Groups](#)
- [Create a Cloud User Account in OpenStack](#)
- [Modify a User Account](#)

Create an OpenStack Project

Projects are the equivalent of tenants or accounts. They function as organizational units in the cloud to which you can assign users.

Prerequisites

Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 Select the admin project from the drop-down menu in the title bar.
- 2 Select **Admin > Identity > Projects**.
- 3 Click **Create Project**.

- 4 Click the **Project Info** tab and configure the project settings.

Setting	Description
Name	Project name.
Description	Optional description of the new project.
Enabled	New projects are enabled by default. Disabling a project prevents cloud users from accessing the project, prevents users from managing launching instances for the project, and can prevent users from logging in if they are assigned only to that project.

- 5 (Optional) Add members to the project by selecting existing cloud users on the **Project Members** tab.
- 6 (Optional) Add groups of members to the project by selecting existing groups of cloud users on the **Project Groups** tab.
- 7 On the **Quota** tab, accept or modify the quota settings.

Quotas are operational limits that you can configure to manage how much system resources are available to a specific project. For example, you can optimize the cloud resources by controlling the number of gigabytes allowed for each tenant. Quotas can be enforced at both the project and user level.

- 8 Click **Create Project** at the bottom of the panel.

The VMware Integrated OpenStack dashboard assigns an ID to the new project, and the project is listed on the Projects page.

Modify a Project

You can update a project to change its name or description, and enable or temporarily disable it.

Important Disabling a project can have negative consequences. For example, if a user is assigned to only that project, they cannot log in to the VMware Integrated OpenStack dashboard. Similarly, the project is not accessible by its members. Project instances continue running, so you must suspend or stop them manually. Project data is retained in case the project is enabled again.

Prerequisites

Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 On the VMware Integrated OpenStack dashboard, select the admin project from the drop-down menu in the title bar.
- 2 Select **Admin > Identity > Projects**.
- 3 Select the project to edit.

- 4 In the Actions column, select **Edit Project** from the drop-down menu.

In the Edit Project dialog box, you can change the project's name and description, and enable and disable it.

- 5 Modify the project settings, and click **Save**.
- 6 (Optional) To change user assignments for a project, on the Projects page, click **Manage Members** for the project to modify.

Option	Action
Assign a user to the current project	Click the plus sign (+) for the user.
Remove a user from the current project,	Click the minus sign (-) for the user.

- 7 Click **Save**.
- 8 To delete one or more projects, return to the Projects page and select the projects to delete.

Note You cannot restore a deleted project.

- a Click **Delete Projects**.
- b At the prompt, confirm the deletion.

Working with Security Groups

A security group is a set of IP filter rules that define networking access and that you can apply to all instances in a project. Group rules are project-specific. Project members can edit the default rules for their group and add new rule sets.

You can use security groups to apply IP rules by creating a new security group with the desired rules or by modifying the rules set in the default security group.

Note A security group can apply either rules or a security policy, but not both.

About the Default Security Group

Each project in VMware Integrated OpenStack has a default security group that is applied to an instance unless another security group is defined and specified. Unless it is modified, the default security group denies all incoming traffic to your instance and permits only outgoing traffic. A common example is to edit the default security group to permit SSH access and ICMP access, so that users can log in to and ping instances.

Create a Security Group

Security groups are sets of IP filter rules that define networking access and are applied to all instances within a project. You can either modify the rules in the default security group or create a security group with custom rules.

To modify an existing rule for a security group, see [Modify the Rules for an Existing Security Group](#)

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Security Groups** tab.
- 5 Click **Create Security Group**.
- 6 Enter a name and description for the new group, and click **Create Security Group**.
The new group appears in the list on the **Security Group** tab.
- 7 Configure the rules for the new group.
 - a Select the new security group and click **Manage Rules**.
 - b Click **Add Rule**.
 - c From the **Rule** drop-down menu, select the rule to add.
The subsequent fields might change depending on the rule you select.
 - d If applicable, specify **Ingress** or **Egress** from the **Direction** drop-down menu.
 - e After you complete the rule definition, click **Add**.
- 8 Configure additional rules if necessary.
- 9 Click the **Access & Security** tab to return to the main page.

Modify the Rules for an Existing Security Group

You can modify a security group by adding and removing rules assigned to that group. Rules define which traffic is allowed to instances that are assigned to the security group.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Security Groups** tab.
- 5 Select the security group to modify and click **Manage Rules**.

- 6 To remove a rule, select the rule and click **Delete Rule**.
- 7 To add a rule, click **Add Rule** and select the custom rule to add from the **Rule** drop-down menu.

Option	Description
Custom TCP Rule	Used to exchange data between systems and for end-user communication.
Custom UDP Rule	Used to exchange data between systems, for example, at the application level.
Custom ICMP Rule	Used by network devices, such as routers, to send error or monitoring messages.
Other Protocol	You can manually configure a rule if the rule protocol is not included in the list.

- a From the **Remote** drop-down list, select **CIDR** or **Security Group**.
- b If applicable, select **Ingress** or **Egress** from the **Direction** drop-down menu.

For TCP and UDP rules, you can open either a single port or a range of ports. Depending on your selection, different fields appear below the Open Port list.

- c Select the kind of access to allow.

Option	Description
CIDR (Classless Inter-Domain Routing)	Limits access only to IP addresses within the specified block.
Security Group	Allows any instance in the specified security group to access any other group instance. You can choose between IPv4 or IPv6 in the Ether Type list.

- 8 Click **Add**.

The new rule appears on the Manage Security Group Rules page for the security group.

Enabling SSH and ICMP Access

You can modify the default security group to enable SSH and ICMP access to instances. The rules in the default security group apply to all instances in the currently selected project.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Security Groups** tab, select the default security group, and click **Manage Rules**.
- 5 Click **Add Rule** and configure the rules to allow SSH access.

Control	Value
Rule	SSH
Remote	CIDR
CIDR	0.0.0.0/0

To accept requests from a particular range of IP addresses, specify the IP address block in the CIDR text box.

Instances will now have SSH port 22 open for requests from any IP address.

6 Click **Add**.

7 From the Manage Security Group Rules page, click **Add Rule** and configure the rules to allow ICMP access.

Control	Value
Rule	All ICMP
Direction	Ingress
Remote	CIDR
CIDR	0.0.0.0/0

8 Click **Add**.

Instances will now accept all incoming ICMP packets.

Use VMware NSX for vSphere Security Policies Through Security Groups

This feature enables the consumption of VMware NSX for vSphere policy from the OpenStack Cloud Management Platform through OpenStack security groups. NSX administrator can define security policies that the OpenStack cloud administrator shares with cloud users. Cloud user can also define their own security groups with rules if the cloud administrator enables regular security groups. This feature can also be used by cloud administrators to insert third-party network services.

Starting with VMware Integrated OpenStack 3.1, Neutron security groups enable administrators to use two new functionalities.

Provider security groups Also known as administrator rules, when configured those security groups are mandatory and apply to all VMs of a given tenant. A provider security group can either be associated with a policy or exist without a policy.

NSX Service Composer - Security Policy security groups For more information, see the *Service Composer* chapter in the *VMware NSX for vSphere Administration Guide*.

Each VMware NSX for vSphere policy can be defined by the OpenStack Cloud Administrator as a default policy by setting the `nsxv_default_policy_id` option in the `custom.yml` file. All new tenants have this policy as their default. More policies can be defined and assigned as mandatory or optional for a given tenant by being associated with either the provider or optional security groups respectively. Tenant users can also create security groups with rules but they cannot override security groups set by the cloud administrator.

After VMware NSX for vSphere policies are enabled, different scenarios can be configured by cloud administrators.

- 1 Cloud administrator can forbid the creation of regular security groups with different options.
 - If only a default security group exists, this default security group is associated with the default policy. Tenant VMs are enforced with the rules defined in the default policy.
 - If the cloud administrator creates a security group with a different policy, tenant VMs can be associated with this security group instead of the default security group and only the rules defined in the current policy are effective.
 - If provider security groups exist, in addition to the policy rules, tenant VMs are also be enforced with the rules defined in the provider security groups.
- 2 Cloud administrator can allow the creation of regular security groups with different options.
 - VMs launched with user-defined regular security groups are only enforced with the rules defined in these security groups.
 - If a provider security group exists, in addition to the rules in the regular security group, tenant VMs are also enforced with the rules defined in the provider security groups. In this case, provider security group rules take precedence over regular security group rules. Similarly, if you use policy-based security groups with regular security groups, policy-based rules take precedence.
 - You can have security groups either with a policy or rules, but not with both.

Manage NSX Service Composer - Security Policy Security Groups Through CLI Commands

Cloud administrators can also change the association of security group policy by using CLI commands through the Integrated OpenStack Manager.

Action	Command Example
Change the associated policy for a security group.	<code>neutron security-group-update --policy=<NSX_Policy_ID> <SECURITY_GROUP_ID></code>
Migrate existing security groups to policy-based security groups by using the <code>nsxadmin</code> utility.	<code>nsxadmin -r security-groups -o migrate-to-policy --property policy-id=<NSX_Policy_ID> --property security-group-id=<SECURITY_GROUP_ID></code>
Note This action deletes existing rules defined by the user. Make sure that you have the appropriate rules in the policy to avoid network disruption.	

Action	Command Example
Enforce provider security groups on existing VM ports	<pre>neutron port-update <PORT_ID> --provider-security-groups list=true <SECURITY_GROUP_ID1> <SECURITY_GROUP_ID2></pre>
Ensure that a new policy, created on the NSX side is placed before all the OpenStack security groups section by using the nsxadmin utility.	<pre>sudo -u neutron nsxadmin --config-file /etc/neutron/neutron.conf --config-file /etc/neutron/plugins/vmware/nsxv.ini -r firewall-sections -o nsx-reorder</pre>
<p>Note When more than one policy-based security groups are enforced on a VM/port, the order in which the policy rules are enforced is controlled by the NSX admin through the firewall section.</p>	

- [Enable VMware NSX for vSphere Security Policies in Neutron](#)

You enable VMware NSX for vSphere security policies in Neutron by modifying the `custom.yml` file.

- [Modify Security Policy in a Security Groups](#)

You can change the security policy that is associated with a security group.

Enable VMware NSX for vSphere Security Policies in Neutron

You enable VMware NSX for vSphere security policies in Neutron by modifying the `custom.yml` file.

Additionally you must set the default security policy for the default security group for a new tenant and optionally, allow or forbid tenants to create own policies.

Procedure

- 1 Log in to the OpenStack Management Server.
- 2 Create `custom.yml` file, if it does not exist.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample
/opt/vmware/vio/custom/custom.yml
```

- 3 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.

- 4 Enable security policies in Neutron using VIO customization by editing the `custom.yml` file according to your configuration.

- a Uncomment and edit `nsxv_use_nsx_policies` value to **true**, set the mandatory default policy for tenants `nsxv_default_policy_id`, and allow or forbid tenants to create their own policies `nsxv_allow_tenant_rules_with_policy: false`, for example:

```
# Configure neutron security groups to use NSX policies
nsxv_use_nsx_policies: true
# (Optional) If use_nsx_policies is true, this policy will be used as the
# default policy for new tenants.
nsxv_default_policy_id: <YOUR_NSX_POLICY_ID>
# (Optional) If use_nsx_policies is True, this value will determine if the
# tenants can add rules to their security groups.
nsxv_allow_tenant_rules_with_policy: false
```

- b Save the `custom.yml` file.

- 5 Push the new configuration to your VMware Integrated OpenStack deployment.

Refresh of the configuration briefly interrupts the OpenStack services.

```
viocli deployment configure
```

Modify Security Policy in a Security Groups

You can change the security policy that is associated with a security group.

You perform the procedure from a VMware Integrated OpenStack controller.

Prerequisites

Procedure

- 1 Log in to the OpenStack Management Server.
- 2 Get a list of the currently defined security groups.

You need the `id` of a security group to see its configuration.

```
neutron-security-group-list
```

- 3 Get the configuration of a security group.

Use the `id` from the previous step.

```
neutron-security-group-show <SECURITY_GROUP_ID>
```

In the output you see the `policy` associated with this security group.

- 4 Change the current policy of a security group with another policy.

```
neutron security-group-update --policy=<NSX_Policy_ID> <SECURITY_GROUP_ID>
```

You have changed the associated security policy to a given security group.

Create a Cloud User Account in OpenStack

Cloud users have a limited set of rights and privileges relative to cloud administrators. Cloud users are limited to the tenants to which they are assigned. Tenants are called projects in OpenStack. Cloud users can create and manage instances, create and manage volumes, create networks, and create new images, among other tasks.

Prerequisites

- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- Verify that a configured OpenStack project is available. See [Create an OpenStack Project](#).

VMware Integrated OpenStack now supports Keystone multi-domain back-end:

- A separate back-end can be on each domain.
- The *local* domain now contains the service users, vioservice user, and admin user. This domain is backed by SQL. The *Default* domain contains either your standard users if using SQL or LDAP users if AD was configured. For convenience, the admin user is also available on the *Default* domain.
- Along with the domain context in Horizon, in the CLI you will also need to specify a domain if not using the Default. OpenStack command lines will always default to the default domain.
- When logging in to the dashboard, users are now prompted for a domain name. To log in successfully, they must enter "default" for the domain name

Procedure

- 1 On the VMware Integrated OpenStack dashboard, select the admin project from the drop-down menu in the title bar.
- 2 Select **Admin > Identity Panel > Users**.
- 3 Click **Create User**.

The Create User dialog appears.

- 4 Ensure that the Domain ID field is set to default and the Domain name is Default.

Users must enter the correct domain name to successfully log into the VMware Integrated OpenStack dashboard.

- 5 Configure the user settings.

Option	Description
User Name	Cloud user name.
Email	Valid email address for the new user.
Password/Confirm Password	Preliminary password for the new user.

Option	Description
Primary Project	Project to which the user is assigned. You cannot create a user account without assigning it to at least one project.
Role	Role to which the user is assigned. A role is a set of rights and privileges. A user assigned that role inherits those rights and privileges.
Enable	To enable the user, check the Enable check box. To enable the user at a later time, leave the Enable check box unchecked.

- 6 Click **Create User** at the bottom of the panel.

The VMware Integrated OpenStack dashboard assigns an ID to the user, and the user now appears on the Users page.

Modify a User Account

As a cloud administrator, you can enable, disable, and delete user accounts, and change passwords for accounts.

Prerequisites

Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 In the VMware Integrated OpenStack dashboard, select the admin project from the drop-down menu in the title bar.
- 2 Select **Identity > Users**.

Option	Action
Enable or disable a user account.	<ol style="list-style-type: none"> a Select the user account to edit. b In the Actions column, click Edit and select Enable User or Disable User from the drop-down list.
Delete one or more user accounts.	<ol style="list-style-type: none"> a Select the user accounts to delete. b Click Delete Users. c At the prompt, confirm the deletion.
Change a password	<ol style="list-style-type: none"> a Select the user account to edit. b In the Actions column, click Edit and then select Change Password. c Change the password as required.

Working with Instances in OpenStack

4

Instances are virtual machines that run in the cloud.

As a cloud administrative user, you can manage instances for users in various projects. You can view, terminate, edit, perform a soft or hard reboot, create a snapshot from, and migrate instances. You can also view the logs for instances or start a VNC console for an instance.

For information about how to use the dashboard to start instances as an end user, see the *VMware Integrated OpenStack User's Guide*.

This section includes the following topics:

- [Import vSphere VMs into VMware Integrated OpenStack](#)
- [Create a Snapshot from an Instance](#)
- [Control the State of an Instance](#)
- [Track Instance Use](#)
- [Use DRS to Control OpenStack Instance Placement](#)
- [Using Affinity and Anti-Affinity to Place OpenStack Instances](#)
- [Apply QoS Resource Allocation to Existing Instances](#)
- [Define Default Nova Storage for OpenStack Instances](#)

Import vSphere VMs into VMware Integrated OpenStack

You can import VMs from vSphere into your VMware Integrated OpenStack deployment and manage them like OpenStack instances.

You import VMs using the Datacenter Command Line Interface (DCLI), which is packaged with the VMware Integrated OpenStack management server, and is powered by the VMware Integrated OpenStack vAPI provider.

Though imported VMs become OpenStack instances, they remain distinct in several ways:

- If the imported VM has multiple disks:
 - Nova snapshot creation is not supported.
 - The Nova resizing operation is not supported.

- Existing networks are imported as provider network type port group, with subnets created with DHCP disabled. This prevents conflict between the DHCP node in OpenStack and the external DHCP server.

Note If the DHCP server cannot maintain the same IP address during lease renewal, the instance information in OpenStack will show the incorrect IP address. For this reason, it is recommended that you use static DHCP bindings on existing DHCP servers. Also, it is not recommended to launch new OpenStack instances on the imported networks since the DHCP address from the external server, if any, might conflict with OpenStack.

- The flavor for the imported VM shows the correct CPU and memory but the root disk incorrectly displays as having 0 GB.

Prerequisites

- Verify that you are running VMware Integrated OpenStack version 4.0.
- Verify that VMware Integrated OpenStack is deployed and running.
- Verify that the VMs to be imported are in the same vCenter.
- Importing VMs is supported with NSX and the VDS plugin for Neutron.

Note If you are running VMware Integrated OpenStack 3.0, you cannot import VMs that are backed by a NSX logical switch. The network backing must be a regular distributed port group. This feature is supported in VMware Integrated OpenStack 3.1 and later.

Procedure

- Add the clusters containing the VMs to be imported to the VMware Integrated OpenStack deployment.
 - In the vSphere Web Client, identify the cluster containing the VMs to be imported.
 - Add the cluster to the VMware Integrated OpenStack deployment as a Nova compute cluster.
 - Repeat for multiple clusters, if necessary.

After the cluster is added as a Nova compute cluster, you can import the VMs.

- Using SSH, log in to the VMware Integrated OpenStack manager.
- Connect to the VMware Integrated OpenStack vAPI endpoint.

The endpoint runs locally.

```
dccli +server http://localhost:9449/api +i
```

This command opens an interactive shell (`dccli`).

4 List all namespaces in the VMware Integrated OpenStack vAPI provider.

```
dcli> com vmware vio
The vio namespace provides namespaces to manage components related to OpenStack and vSphere
Available Namespaces:
vm
```

5 (Optional) List the commands related to importing unmanaged VMs.

Unmanaged VMs are VMs in VMware Integrated OpenStack that are not managed as OpenStack instances. In this case, the unmanaged VMs include the VMs in the cluster you added to the Compute node.

```
dcli> com vmware vio vm unmanaged
The Unmanaged namespace provides commands to manage virtual machine not under OpenStack
Available Commands:
importall  Imports all unmanaged virtual machines into OpenStack
importvm   Imports given virtual machine into OpenStack
list      Enumerates the list of unmanaged virtual machines
```

6 (Optional) List all unmanaged VMs in a specific target cluster that you added to the Nova Compute node.

```
com vmware vio vm unmanaged list --cluster <vcenter cluster mor-id>
```


7 Import VMs into VMware Integrated OpenStack.

You can import all VMs or a specific VM.

a To import all VMs:

```
com vmware vio vm unmanaged importall [-h] --cluster CLUSTER [--tenant-mapping
{FOLDER,RESOURCE_POOL}] [--root-folder ROOT_FOLDER]
                                [--root-resource-pool ROOT_RESOURCE_POOL]
```

Option	Description
<code>--cluster CLUSTER</code>	Specify the Nova compute cluster where the VMs are located.
<code>--tenant-mapping {FOLDER,RESOURCE_POOL}</code>	Specify whether to map the vSphere VMs to OpenStack projects based on their location in folders or resource pools. This parameter is optional. If no tenant mapping is specified, imported VMs become instances in the import_service project in OpenStack.
<code>--root-folder ROOT_FOLDER</code>	Optionally, if you specified FOLDER for the <code>tenant-mapping</code> parameter, you can specify the name of the root folder containing the VMs to be imported. <ul style="list-style-type: none"> All the VMs in the specified root folder are imported, including those contained in sub-folders. The VMs will be imported as instances into an OpenStack project with the same name as the specified root folder. If the root folder contains VMs in sub-folders, those VMs will be imported into OpenStack projects with the same names as the sub-folders. <p>Note If no root folder is specified, the name of the top level folder in the cluster is used by default.</p>
<code>--root-resource-pool ROOT_RESOURCE_POOL</code>	Optionally, if you specified RESOURCE_POOL for the <code>tenant-mapping</code> parameter, you can specify the name of the root resource pool containing the VMs to be imported. <ul style="list-style-type: none"> All the VMs in the specified root resource pool are imported, including those contained in child resource pools. The VMs will be imported as instances into an OpenStack project with the same name as the specified root resource pool.

Option	Description
	<ul style="list-style-type: none"> If the root resource pool contains VMs in child resource pools, those VMs will be imported into OpenStack projects with the same names as the child resource pools.

- b To import a specific VM:

```
com vmware vio vm unmanaged importvm [-h] \
  --vm VM [--tenant TENANT] [--nic-mac-address NIC_MAC_ADDRESS] \
  [--nic-ipv4-address NIC_IPV4_ADDRESS]
```

Option	Description
<code>--vm VM</code>	Specify the <code>vm-<id></code> of the specific VM to be imported. You can view the ID values of all VMs to be imported by running the <code>com vmware vio vm unmanaged list</code> command.
<code>--tenant TENANT</code>	Specify the OpenStack project where the imported VM will reside as an OpenStack instance. This parameter is optional. If unspecified, imported VMs become instances in the <code>import_service</code> project in OpenStack.
<code>--nic-mac-address NIC_MAC_ADDRESS</code>	Optionally, provide the MAC address for the VM's NIC. If the import process is unable to discover this value, the import will fail. This parameter enables you manually enter the NIC MAC address. Note If specified, you must also provide the <code>nic-ipv4-address</code> parameter.
<code>--nic-ipv4-address NIC_IPV4_ADDRESS</code>	Optionally, provide the IP address for the VM's NIC. If the import process is unable to discover this value, the import will fail. This parameter enables you manually enter the NIC IP address. Note If specified, you must also provide the <code>nic-mac-address</code> parameter.

- 8 (Optional) You can enable or disable the relocation and renaming of imported VMs by modifying the `custom.yml` file.

This option is enabled by default.

- a If you have not already done so, implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- b To disable the relocation and renaming of imported VMs, uncomment the following parameter in the `custom.yml` file.

```
nova_import_vm_relocate: false
```

- c Save the `custom.yml` file.

Create a Snapshot from an Instance

With snapshots, you can create new images from running instances.

You can create a snapshot of an instance directly from the Instances page.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Instances**.
- 4 In the Actions column, click **Create Snapshot**.

The snapshot appears on the Images & Snapshots page.

Control the State of an Instance

As a cloud administrative user, you can pause, unpaue, suspend, resume, soft or hard reboot, or terminate an instance.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Instances**.
- 4 Select the instance whose state you want to manage.
- 5 In the Actions column, click **More** and select the state from the drop-down menu.

Items that appear in red text are disabled.

Track Instance Use

You can track the use of instances for each project. You can track costs per month by showing metrics like the number of VCPUs, disks, RAM, and uptime of all of your instances.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Overview**.

The Overview page shows the usage summary and project-specific usage information. You can specify a period of time for the usage information. Optionally, you can download a CSV summary.

- 4 (Optional) Specify a period of time for reporting and click **Submit**.
- 5 (Optional) Click **Download CSV Summary** to download a report of the usage.

Use DRS to Control OpenStack Instance Placement

As a cloud administrator, you can use vSphere DRS settings to control how specific OpenStack instances are placed on hosts in the Compute cluster. In addition to the DRS configuration, you also modify the metadata of source images in OpenStack to ensure that instances generated from those images are correctly identified for placement.

Prerequisites

- Verify that you are running VMware Integrated OpenStack version 2.0.x or later.
- Verify that VMware Integrated OpenStack is running in vSphere. Go to **Home > VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name]**.
- At least one dummy VM in the Compute cluster to use as a template to create a DRS VM group.

Procedure

1 [Define VM and Host Groups for Placing OpenStack Instances](#)

In the vSphere Web Client, create VM and host groups to contain and manage specific OpenStack instances.

2 [Create a DRS Rule for OpenStack Instance Placement](#)

In the vSphere Web Client, create a DRS rule to manage the distribution of OpenStack instances in a VM group to a specific host group.

3 [Apply VM Group Settings to Image Metadata](#)

You can modify the metadata of a source image to automatically place instances into VM groups. VM groups are configured in the vSphere Web Client and can be further used to apply DRS rules.

Define VM and Host Groups for Placing OpenStack Instances

In the vSphere Web Client, create VM and host groups to contain and manage specific OpenStack instances.

Procedure

- 1 Log in to the vSphere Web Client.
- 2 Go to the vCenter Hosts and Clusters view.
- 3 Select the Compute cluster configured for the VMware Integrated OpenStack deployment.
- 4 Click the **Manage** tab.
- 5 Click **Settings** and click **vSphere DRS**.

- 6 Verify the following settings configuration:
 - **DRS** is enabled.
 - **DRS Automation** is set to Fully Automated or Partially Automated .
 - **Power Management** is set to Off.
- 7 Click **VM/Host Groups**.
- 8 Create a VM group.
 - a Click **Add**.
 - b Enter a name for the new VM group.
 - c From the **Type** drop-down menu, select **VM Group**.
 - d Click **Add**.
 - e On the **Filter** tab, select the dummy VM to create an empty VM group.
You created the dummy VM in an earlier task in this sequence.
 - f Click **OK**.
- 9 Create a Host group.
 - a Click **Add**.
 - b Enter a name for the new host group.
 - c From the **Type** drop-down menu, select **Host Group**.
 - d Click **Add**.
 - e On the **Filter** tab, add members to the group by selecting one or more hosts.
 - f Click **OK**.

Both groups now appear in the VM/Host Groups list on the VM/Host page.

What to do next

You can now create a rule that determines how OpenStack instances assigned to the VM group are distributed on the hosts in the host group. See [Create a DRS Rule for OpenStack Instance Placement](#).

Create a DRS Rule for OpenStack Instance Placement

In the vSphere Web Client, create a DRS rule to manage the distribution of OpenStack instances in a VM group to a specific host group.

If you are continuing from [Define VM and Host Groups for Placing OpenStack Instances](#), skip ahead to [Step 5](#) .

Prerequisites

- Define at least one VM group .

- Define at least one host group.

See [Define VM and Host Groups for Placing OpenStack Instances](#).

Procedure

- 1 Log in to the vSphere Web Client.
- 2 Go to the vCenter Hosts and Clusters view and select the Compute cluster configured for the VMware Integrated OpenStack deployment.
- 3 Click the **Manage** tab, and go to **Settings > vSphere DRS**.
- 4 Verify the following settings configuration:
 - **DRS** is enabled .
 - **DRS Automation** is set to Fully Automated or Partially Automated .
 - **Power Management** is set to Off.
- 5 Click **VM/Host Rules**.
- 6 Click **Add**.
- 7 Enter a name for the new rule and select or deselect the **Enable rule** option to enable or disable the rule.
- 8 From the **Type** drop-down menu, select **Virtual Machines to Hosts**.
- 9 From the **VM Group** drop-down menu, select the VM group that identifies the OpenStack instances you want to place .
- 10 Select the **Must run on hosts in group** specification. .
- 11 Select a specification for the rule.

Setting	Description
Must run on hosts in group	OpenStack instances in the specified VM group must run on hosts in the specified host group.
Should run on hosts in group	OpenStack instances in the specified VM group should, but are not required, to run on hosts in the specified host group.
Must not run on hosts in group	OpenStack instances in the specified VM group must never run on host in the specified host group.
Should not run on hosts in group	OpenStack instances in the specified VM group should not, but may, run on hosts in the specified host group.

- 12 From the **Host Group** drop-down menu, select the host group that contains the hosts on which the OpenStack instances will be placed .
- 13 Click **OK**.

The rule now determines that OpenStack instances in the specified VM group must run on hosts in the specified host group.

What to do next

In the VMware Integrated OpenStack dashboard, you can now modify the metadata for a specific image to ensure that all instances generated from that image are automatically included in the VM group, and therefore subject to the DRS rule.

Apply VM Group Settings to Image Metadata

You can modify the metadata of a source image to automatically place instances into VM groups. VM groups are configured in the vSphere Web Client and can be further used to apply DRS rules.

Prerequisites

- Verify that a VM group is configured in the vSphere Web Client for the Compute cluster.
- Verify that the DRS VM group name is defined in the vSphere Web Client. See [Use DRS to Control OpenStack Instance Placement](#).

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Images**.
- 4 Click the image to modify.
- 5 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
- 6 Add the DRS VM group metadata property to the image metadata.

The Update Metadata dialog box displays two columns. The column on the right displays metadata tags already applied to the image, and the one on the left displays available metadata tags, which are grouped by categories, such as Guest Customization, Instance Config Data, and so on.

- a In the Available Metadata column, select the **VMware Driver Options > DRS VM group** property.
- b Click the plus sign (+) to add the property to the image metadata.

The **vmware_vm_group** metadata property is highlighted in the Existing Metadata column.

- c For the metadata value, enter DRS VM group name as defined in the vSphere Web Client.
- d To remove a metadata tag from the image definition, click the minus sign (-).

- 7 Click **Save**.

All instances generated from this source image are automatically assigned to the specified VM group in the VMware Integrated OpenStack deployment in vCenter.

Using Affinity and Anti-Affinity to Place OpenStack Instances

The Nova scheduler provides filters that you can use to ensure that OpenStack instances are automatically placed on the same host (affinity) or separate hosts (anti-affinity).

You apply the affinity or anti-affinity filter as a policy to a server group. All instances that are members of the same group are subject to the same filters. When you create an OpenStack instance, you can specify the server group to which the instance will belong and therefore what filter will be applied.

You can perform this configuration using either the OpenStack CLI or ServerGroup API. You cannot perform this configuration in the VMware Integrated OpenStack Horizon dashboard.

This approach to placing OpenStack instances is tenant-based. Affinity and anti-affinity determine the relationship among instances in the same server group, but they cannot determine the hosts on which the instances are placed in vCenter. For an administrator-based approach that provides greater control, see [Use DRS to Control OpenStack Instance Placement](#).

Create Instances with an Affinity or Anti-Affinity Policy by Using the CLI

You can place instances using affinity or anti-affinity by creating a server group in OpenStack and applying desired filter as a group policy. All instances that are members of the server group will be subject to the affinity or anti-affinity policy. You can perform this configuration using the CLI.

Prerequisites

- Verify that the intended filter configuration does not conflict with any existing administrative configuration, such as DRS rules that manage instance placement on hosts.
- Verify that you are running VMware Integrated OpenStack version 2.0.x or later.
- Verify that VMware Integrated OpenStack is running.
- Verify that you are using a Python nova-client version 2.17.0.6 or later as required for the ServerGroup API. Go to http://docs.openstack.org/user-guide/common/cli_install_openstack_command_line_clients.html.

Procedure

- 1 Using SSH, log in to the nova-client.
- 2 (Optional) Obtain the ID of the image you will use to create the instance.
You can use the `nova image-list` command to view the list of available images and their ID values.
- 3 (Optional) Obtain the ID of the flavor you will use to define the instance .
You can use the `nova flavor-list` command to view the list of flavor definitions and their ID values.

4 Create a new server group with the intended policy.

a Create a server group with the affinity policy:

```
nova server-group-create GROUP_NAME affinity
```

b Create a server group with the anti-affinity policy:

```
nova server-group-create GROUP_NAME anti-affinity
```

In both case, the CLI returns the auto-generated server group UUID, name, and policy.

5 Launch a new instance, using the `--image`, `--flavor`, and `--hint` flags to apply the server group affinity policy .

```
nova boot --image IMAGE_ID --flavor FLAVOR_ID --hint group=SERVER_GROUP_UUID INSTANCE_NAME
```

6 (Optional) Confirm that the new rule and the server group instances appear and are running correctly in the VMware Integrated OpenStack deployment in vCenter Server.

The details appear in the **Manage > Settings > VM/Host Rules** page for the Compute cluster.

Create Instances with an Affinity or Anti-Affinity Policy Using the API

You can place instances using affinity or anti-affinity by creating a server group in OpenStack and applying desired filter as a group policy. All instances that are members of the server group will be subject to the affinity or anti-affinity policy. You can perform this configuration using the ServerGroup API from the Python nova-client.

Prerequisites

- Verify that the intended anti-affinity filter configuration does not conflict with any existing administrative configuration, such as DRS rules that manage instance placement on hosts.
- Verify that you are running VMware Integrated OpenStack version 2.0.x or later.
- Verify that VMware Integrated OpenStack is running.
- Verify that you are using a Python nova-client version 2.17.0.6 or later, as required for the ServerGroup API. Go to http://docs.openstack.org/user-guide/common/cli_install_openstack_command_line_clients.html.

Procedure

- 1 Create a new server group with an anti-affinity policy.

```
POST /v2/TENANT_ID/os-server-groups
```

```
{
  "server_group": {
    "name": "SERVER_GROUP_NAME",
    "policies": ["POLICY_TYPE"]
  }
}
```

Option	Description
TENANT_ID	ID value for the OpenStack tenant.
SERVER_GROUP_NAME	Specify the name for the server group.
POLICY_TYPE	Specify either affinity or anti-affinity .

- 2 Launch a new instance, including the `os:scheduler_hints` argument with the server group ID in the `GET /servers` command.

```
... "os:scheduler_hints": {"group": "SERVER_GROUP_UUID"}
```

- 3 (Optional) Confirm that the new rule and the server group instances appear and are running correctly in the VMware Integrated OpenStack deployment in vCenter.

The rule details appear in the **Manage > Settings > VM/Host Rules** page for the Compute cluster.

Apply QoS Resource Allocation to Existing Instances

You can apply QoS resource allocation settings to an existing instance by resizing the instance in the VMware Integrated OpenStack dashboard.

Prerequisites

- Requires an OpenStack flavor with the desired QoS resource allocation settings. See [Configure QoS Resource Allocation for Instances Using Flavor Metadata](#).
- Requires VMware Integrated OpenStack version 2.0.x or greater.
- Verify that VMware Integrated OpenStack is running in vSphere.
- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.

- 3 Select **Admin > System > Instances**.
- 4 Click the hyperlinked name of the instance to access the Instance Details page.
- 5 Click the down arrow (next to the **Create Snapshot** button) and choose **Resize Instance**.
- 6 In the **Flavor Choice** tab, open the **New Flavor** drop-down list and select the flavor with the desired QoS resource allocations
- 7 Click **Resize**.

The resizing process may take a few minutes.

The instance is now subject to the QoS settings as defined in the flavor metadata.

Define Default Nova Storage for OpenStack Instances

To ensure that OpenStack instances booted from a volume use the correct volume type, you can create and apply policy-based management settings, which are called PBM policies.

After you enable the storage policy in the `custom.yml` file, you apply the policy by modifying the metadata of an OpenStack flavor. All instances created by using that flavor inherit the storage policy configuration.

Procedure

- 1 Implement the `custom.yml` file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Edit the `/opt/vmware/vio/custom/custom.yml` file by uncommenting the PBM options.

```
#####
# PBM options
#####

# (string) The PBM default policy to use when no policy is associated with a flavor (Mandatory) if
nova_pbm_enabled is set to True.
nova_pbm_default_policy: nova

# (boolean) The PBM status. Set this to True to enable storage policies for nova flavors.
nova_pbm_enabled: False
```

- 3 Set the `nova_pbm_enabled` parameter to **True**.

```
nova_pbm_enabled: True
```

- 4 Save the `custom.yml` file.

- 5 Apply the policy to an OpenStack flavor as metadata.
 - a Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
 - b Select the admin project from the drop-down menu in the title bar.
 - c Select **Admin > System > Flavors**.
 - d (Optional) Create a flavor specific to the intended use of this metadata property.
Create a custom flavor to contain the specific configuration. This action leaves the original flavor configuration intact and available for other instance creation.
 - e Select the flavor to modify.
 - f In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
 - g Enter **vmware:storage_policy** in the **Custom** field.
 - h Click the plus sign (+) next to the **Custom** field.
In the column under Existing Metadata, the newly added metadata property appears.
 - i Enter **nova** as the metadata property value.
- 6 Click **Save**.

The default Nova storage policy is applied to all future OpenStack instances that are created from this flavor.

Working with Volumes and Volume Types in OpenStack

5

Volumes are block storage devices that you attach to instances to enable persistent storage.

As a cloud administrative user, you can manage volumes and volume types for users in various projects. You can create and delete volume types, and you can view and delete volumes.

Cloud users can attach a volume to a running instance or detach a volume and attach it to another instance at any time. For information about how to use the dashboard to create and manage volumes as an end user, see the *VMware Integrated OpenStack User Guide*.

This section includes the following topics:

- [Modify the Default Cinder Volume Adapter Type](#)
- [Create a Volume Type](#)
- [Delete a Volume Type](#)
- [Migrating Volumes Between Datastores](#)

Modify the Default Cinder Volume Adapter Type

Starting with VMware Integrated OpenStack 3.1, you can change the default adapter type for newly created volumes by changing `vmware_adapter_type` parameter using a `custom.yml` file.

By default, empty volumes are always created and attached to a LsiLogic controller. When a volume is created from image, Cinder respects the `vmware_adapter_type` property of the image and creates the corresponding controller. For newly created volumes you set the adapter type by using the `cinder_volume_default_adapter_type` parameter in the `custom.yml` file with one of the following values.

Value	Description
lsiLogic	Sets the default adapter type to LSI Logic
busLogic	Sets the default adapter type to Bus Logic
lsiLogicsas	Sets the default adapter type to LSI Logic SAS
paraVirtual	Sets the default adapter type to VMware Paravirtual SCSI
ide	Sets the default adapter type to IDE

Procedure

- 1 Implement the custom.yml file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Open the /opt/vmware/vio/custom/custom.yml file in a text editor.
 - a Uncomment the cinder_volume_default_adapter_type parameter.
 - b Change the setting with a custom value, for example **lsiLogicsas**.

```
#####
# cinder-volume options
#####

# Default volume adapter type; valid values are 'lsiLogic',
# 'busLogic', 'lsiLogicsas', 'paraVirtual' and 'ide'. (string value)
#cinder_volume_default_adapter_type: 'lsiLogicsas'
```

- 3 Save the custom.yml file.
- 4 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment configure
```

Note Pushing the configuration briefly interrupts OpenStack services.

Create a Volume Type

If you have cloud administrative permissions, you can manage block storage volumes and volume types for users. After you create a volume type, you use a CLI command to associate it with an existing vCenter storage-based policy. The storage policy defines one or more datastores for the volume type to use.

Prerequisites

- Verify that the storage policy to be associated with the volume type exists. See the [vSphere product documentation](#).
- Verify the name of the storage policy. This value is required when you run the CLI command to associate the volume type with the storage policy.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **System Panel > Volumes**.

The Volumes page lists the volumes that are configured and available to the current user.

- 4 Click the **Volume Types** tab.
- 5 Click **Create Volume Type**.
- 6 Enter a name for the volume type.
- 7 Enter a description for the volume type and then click **Create Volume Type**.
- 8 Associate the volume type with a storage policy.
 - a Log in to one of the controllers in VMware Integrated OpenStack.
 - b Run the cinder command to associate the volume type with a storage policy.

```
cinder type-key name-of-volume-type set vmware:storage_profile=name-of-storage-profile
```

This example uses the following parameters and settings.

Parameter or Setting	Description
<i>name-of-volume-type</i>	Name of the volume type that you defined when you created the volume type.
<i>vmware:storage_profile=name-of-storage-profile</i>	Assigns storage policy by the name defined in vSphere.

- 9 (Optional) If you want to override the default adapter type, associate the volume type with another adapter type.

```
cinder type-key name-of-volume-type set vmware:adapter_type=name-of-adapter-type
```

For adapter type, you can select between one of the following values.

Value	Description
lsiLogic	Sets the adapter type to LSI Logic
busLogic	Sets the adapter type to Bus Logic
lsiLogicsas	Sets the adapter type to LSI Logic SAS
paraVirtual	Sets the adapter type to VMware Paravirtual SCSI
ide	Sets the adapter type to IDE

Delete a Volume Type

As a cloud administrative user, you can manage volumes and volume types for users in projects.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Volumes**.

The Volumes page lists the volumes that are currently configured and available to the current user.

- 4 Select the volume types to delete.

- 5 Click **Delete Volume Types**.
- 6 At the prompt, confirm the deletion.

Migrating Volumes Between Datastores

You can safely migrate Cinder volumes between datastores. This enables you to replace datastores, increase resources and capacity, and preserve volumes without taking them offline. The process for migrating volumes depends on several factors. For example, the process is very straightforward if the volume is not attached to an instance. If a volume is attached to an instance, you must migrate the instance.

Note You cannot migrate any volume that has snapshots attached. You must first detach the snapshots.

Migrate All Volumes from a Specified Datastore

You can quickly evacuate all volumes from a specified datastore, automatically migrating them to other datastores in the same datastore cluster.

Prerequisites

- Verify that the specified datastore is part of a datastore cluster.
- Verify that Storage DRS is enabled in `Not Automation (Manual Mode)` for the datastore cluster.
- Verify that the volume does not have any snapshots attached. If so, you must detach them first.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 Switch to root user.

```
sudo su -
```

- 3 Prepare the volume for migration.

This step prepares all volumes on the specified datastore for migration.

```
viocli ds-migrate-prep [--d DEPLOYMENT] DC_NAME DS_NAME
```

Option	Description
<code>--d DEPLOYMENT</code>	Indicates the name of the VMware Integrated OpenStack deployment.
<code>DC_NAME</code>	Indicates the data center name.
<code>DS_NAME</code>	Indicates the datastore name.

- 4 Place the datastore in maintenance mode.

See the [vSphere product documentation](#).

When you place the datastore in maintenance mode, the datastore is evacuated and the volumes automatically migrate to other datastores in the same datastore cluster.

Migrate Unattached Cinder Volumes

You can migrate Cinder volumes that are unattached to instances to specified target datastores.

Prerequisites

Verify that the volume does not have any snapshots attached. If so, you must detach them first.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 Switch to root user.

```
sudo su -
```

- 3 Migrate the volume.

```
viocli volume-migrate [-d [NAME]] \  
                    [--source-dc [SRC_DC_NAME]] [--source-ds [SRC_DS_NAME]] \  
                    [--volume-ids [VOLUME_UUIDS]] [--ignore-storage-policy] \  
                    DEST_DC_NAME DEST_DS_NAME [-h] [-v]
```

Parameter	Mandatory or Optional	Description
<code>-d, --deployment NAME</code>	Automatic	Name of the deployment in which the volumes to be migrated. Applied automatically. The default value is the name of the current deployment.
<code>--source-dc SRC_DC_NAME</code>	Mandatory unless VOLUME_UUIDS is specified.	Identifies the source data center. Used with the <code>--source-ds</code> parameter uniquely to identify the datastore.
<code>--source-ds SRC_DS_NAME</code>	Mandatory unless VOLUME_UUIDS is specified.	Used with the <code>--source-dc</code> parameter uniquely to identify the datastore. For example, the following command migrates all the volumes from datastore DS-01 in data center DC-01 to datastore DS-02 in data center DC-02. <code>viocli volume-migrate --source-dc DC-01 --source-ds DS-01 DC-02 DS-02</code>
<code>--volume-ids VOLUME_UUIDS</code>	Mandatory unless SRC_DC_NAME and SRC_DS_NAME are specified.	Migrates one or more individual volumes specified by UUID value. To specify more than one volume, separate the UUIDs by commas. For example, the following command migrates two volumes specified by their UUID values to datastore DS-01 in data center DC-01. <code>viocli volume-migrate --volume-ids 25e121d9-1153-4d15-92f8-c92c10b4987f, 4f1120e1-9ed4-421a-b65b-908ab1c6bc50 DC-01 DS-01</code>

Parameter	Mandatory or Optional	Description
<code>--ignore-storage-policy</code>	Optional	Ignores storage policy compliance check. Include this parameter to prevent migration failure if the migrated volume includes a storage policy with which the destination datastore does not comply.
<code>DEST_DC_NAME</code>	Mandatory	Specifies the destination data center.
<code>DEST_DS_NAME</code>	Mandatory	Specifies the destination datastore.
<code>-h, --help</code>	Optional	Show the use and arguments for this command.
<code>-v, --verbose</code>	Optional	Enter verbose mode.

Migrate Attached Cinder Volumes

To migrate an attached cinder volume to a different datastore, you must migrate the virtual machine that corresponds to the instance to which it is attached.

Prerequisites

Verify that the volume does not have any snapshots attached. If so, you must detach them first.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 Switch to root user.

```
sudo su -
```

- 3 Prepare the volume for migration.

This step prepares all volumes on the specified datastore for migration.

```
viocli ds-migrate-prep [-d DEPLOYMENT] DC_NAME DS_NAME
```

Option	Description
<code>-d DEPLOYMENT</code>	Indicates the name of the VMware Integrated OpenStack deployment.
<code>DC_NAME</code>	Indicates the data center name.
<code>DS_NAME</code>	Indicates the datastore name.

- 4 Log in to the vSphere Web Client.
- 5 Locate the virtual machine that corresponds to the Nova instance to which the volume is attached.
- 6 Use the Storage vMotion feature in the vSphere Web Client to migrate the virtual machine to a different datastore.

The volume migrates to the new datastore, but only the disk of the shadow VM moves to the new datastore. The shadow VM remains on the old datastore with no disk.

See the [vSphere product documentation](#) about using Storage vMotion.

- 7 (Optional) To fix the disk of the shadow VM, run a volume detach procedure.

The detach operation disconnects the volume from the instance. Failures to read or write from the volume might occur.

Managing Images for the Image Service

6

In the OpenStack context, an image is a file that contains a virtual disk from which you can install an operating system on a VM. You create an instance in your OpenStack cloud by using one of the images available. The VMware Integrated OpenStack Image Service component natively supports images that are packaged in the ISO, OVA, and VMDK formats.

If you have existing images in vSphere that you want to use in OpenStack, you can export them in one of the supported formats and upload them to the Image Service. If you obtain an image that is in an unsupported format, you can convert it as part of the import process. Unsupported formats are RAW, QCOW2, VDI, and VHD.

This section includes the following topics:

- [Import Images to the Image Service](#)
- [Modify Image Settings](#)
- [Modify Image Resource Metadata](#)
- [Configuring Images for Windows Guest Customization](#)
- [Configure QoS Resource Allocation for Instances Using Image Metadata](#)
- [Delete an Existing Image](#)
- [Migrating Images](#)
- [Add a VM Template as an Image](#)
- [Configure Images to Enable Live Resize of VMs Deployed From That Image](#)
- [Modify the Default Behavior for Nova Snapshots](#)
- [Modify the Default Cinder Upload-to-Image Behavior](#)

Import Images to the Image Service

You can use CLI commands or the VMware Integrated OpenStack dashboard to import images.

Prerequisites

To be successfully imported, verify that the image is in one of the natively supported image formats (ISO, OVA, VMDK) or in a format that can be converted during the import process (RAW, QCOW2, VDI, VHD).

Procedure

1 Import Images Using the Horizon Dashboard

You can import images directly in the VMware Integrated OpenStack Horizon dashboard.

2 Import Images in Supported Formats Using the CLI

You can make images available for use in instances by importing images to the Image Service datastore .

3 Import Images in Unsupported Formats by Using the CLI

You can import images in unsupported image formats such as RAW, QCOW2, VDI, or VHD using the `glance-import` tool in the CLI. This tool automatically converts the source image to the VMDK format.

Import Images Using the Horizon Dashboard

You can import images directly in the VMware Integrated OpenStack Horizon dashboard.

Prerequisites

- Verify that the image is packaged in the ISO, VMDK, OVA, RAW, QCOW2, VDI, or VHD format.
- If the source image format is RAW, QCOW2, VDI, or VHD, verify that the source image is hosted on a server without credentials to allow plain HTTP requests.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Images**.
- 4 On the Images page, click **Create Image**.
- 5 Configure the image.

Option	Action
Name	Enter a name for the new image.
Description	(Optional) Enter a description for the new image.
Image Source	Select the image source. If the source image format is RAW, QCOW2, VDI, or VHD, you must select the Image Location option.
Disk Format	Select the disk format.
Disk Type	Select the disk type. Images in the RAW, QCOW2, VDI, and VHD formats are automatically introspected to capture their properties and converted to the VMDK format during the import process.
Adapter Type	Select the adapter type.
Architecture	Accept the default.
OS Type	Select the type of operating system.

Option	Action
Minimum Disk (GB)	Specify the minimum disk size for the image in GB.
Minimum RAM (GB)	Specify the minimum RAM for the image.
Public	Select to make the image visible and available to all tenants.
Protected	Select to prevent the image from being deleted.

6 Click **Create Image**.

The Images page now includes the newly added image.

The image is now ready for deployment in OpenStack instances.

Import Images in Supported Formats Using the CLI

You can make images available for use in instances by importing images to the Image Service datastore .

To import an image in a non-supported format such as RAW, QCOW2, VDI, or VHD, see [Import Images in Unsupported Formats by Using the CLI](#).

Prerequisites

- Verify that you configured one or more Image Service datastores.
- Obtain the image, for example, `ubuntuLTS-sparse.vmdk`.
- Verify that the images are packaged in the ISO, VMDK, or OVA format.

Procedure

- 1 Log in to the OpenStack management cluster as a user with administrative privileges to upload the image to the Image Service component.
- 2 Run the `glance image-create` command to obtain, define, and import the image.

```
glance --os-auth-token $token --os-image-url http://123.456.7.8:9292 \
  image-create name="ubuntu-sparse" \
  disk_format=vmdk \
  container_format=bare \
  --visibility="public" \
  --property vmware_adapertype="lsiLogicsas" \
  --property vmware_disktype="sparse" \
  --property vmware_ostype="ubuntu64Guest" < ubuntuLTS-sparse.vmdk
```

This example uses the following parameters and settings.

Parameter or Setting	Description
<code>--os-image-url</code> <code>http://123.456.7.8:9292</code>	The URL of the source image.
<code>name="ubuntu-sparse"</code>	The name of the source image, in this case, ubuntu-sparse .
<code>disk_format=vmdk</code>	The disk format of the source image. You can specify ISO, VMDK, or OVA.

Parameter or Setting	Description						
<code>container_format=bare</code>	The container format indicates if the image is in a format that contains metadata about the actual virtual machine. Because the container format string is not currently used by Glance, it is recommended to specify bare for this parameter.						
<code>--visibility="public"</code>	The privacy setting for the image in OpenStack. When set to public , the image is available to all users. When set to private , the image is available only to the current user.						
<code>--property vmware_adaptype="lsiLogicsas"</code>	<p>During import, the VMDK disk is introspected to capture its adapter type property.</p> <p>You also have the option of using the <code>vmware_adaptype</code> to specify adapter type.</p> <p>Note If you are using a disk with the paraVirtual or LSI Logic SAS adapter type, it is recommend that you use this parameter. For example, <code>vmware_adaptype= lsiLogicsas</code> or <code>vmware_adaptype= paraVirtual</code>.</p>						
<code>--property vmware_disktype="sparse"</code>	<p>During import, the VMDK disk type is introspected to capture its disk type property.</p> <p>You also have the option of specifying disk type using the <code>vmware_disktype</code> property.</p> <table border="0"> <tr> <td>sparse</td> <td>This disktype property applies to monolithic sparse disks.</td> </tr> <tr> <td>preallocated</td> <td>This disktype property applies to VMFS flat disks, including thick, zeroedthick, or eagerzeroedthick. This is the default property if none is specified.</td> </tr> <tr> <td>streamOptimized</td> <td>This disktype property applies to Monolithic Sparse disks, optimized for streaming. You can convert disks dynamically to and from this format with minimal computational costs.</td> </tr> </table>	sparse	This disktype property applies to monolithic sparse disks.	preallocated	This disktype property applies to VMFS flat disks, including thick, zeroedthick, or eagerzeroedthick. This is the default property if none is specified.	streamOptimized	This disktype property applies to Monolithic Sparse disks, optimized for streaming. You can convert disks dynamically to and from this format with minimal computational costs.
sparse	This disktype property applies to monolithic sparse disks.						
preallocated	This disktype property applies to VMFS flat disks, including thick, zeroedthick, or eagerzeroedthick. This is the default property if none is specified.						
streamOptimized	This disktype property applies to Monolithic Sparse disks, optimized for streaming. You can convert disks dynamically to and from this format with minimal computational costs.						
<code>--property vmware_ostype="ubuntu64Guest"</code>	The name of the image file after it is imported to the Image Service. In the example above, the resulting name will be <code>ubuntuLTS-sparse.vmdk</code> .						

3 (Optional) In the Compute component, confirm that the image was successfully imported.

```
$ glance image-list
```

The command returns a list of all images that are available in the Image Service.

Import Images in Unsupported Formats by Using the CLI

You can import images in unsupported image formats such as RAW, QCOW2, VDI, or VHD using the `glance-import` tool in the CLI. This tool automatically converts the source image to the VMDK format.

You can also use the `glance-import` tool to import images in the supported OVA and VMDK formats.

Prerequisites

- Verify that the image is packaged in the RAW, QCOW2, VDI, or VHD format.
- To allow plain HTTP requests, verify that the image is hosted on a server without credentials.
- Verify that the VMware Integrated OpenStack controller can access the hosted server where the image is stored.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.
- 2 From the VMware Integrated OpenStack manager, use SSH to log in to the controller01 node.
- 3 Switch to root user.

```
sudo su -
```

- 4 Execute the `cloudadmin.rc` file.

```
source cloudadmin.rc
```

- 5 Configure the controller01 node to use the internal VIP.

```
export OS_AUTH_URL=http://INTERNAL_VIP:35357/v2.0
```

- 6 To import the image, run the `glance-import` command.

```
glance-import import --name image_name --url image_http_url --image-format supported_image_format
```

Parameter	Description
image-name	Specify the name for the image as it will appear in the Image Service.
image_format	Specify the format of the source image file. Non-VMDK images are converted automatically to the VMDK format. The following formats are supported: <ul style="list-style-type: none"> ■ VMDK ■ OVA ■ RAW ■ QCOW2 ■ VDI ■ VHD
image_http-url	Provide the HTTP location of the source image file.

For example:

```
glance-import cirros-img qcow2 https://launchpad.net/cirros/trunk/0.3.0/+download/cirros-0.3.0-x86_64-disk.img
```


The CLI displays the task information and status, including the task ID and image ID.

```
Created import task with id 5cdc4a04-5c68-4b91-ac44-37da07ec82ec
Waiting for Task 5cdc4a04-5c68-4b91-ac44-37da07ec82ec to finish.
Current Status.. SUCCESS
Image cirros-img created with ID: 2120de75-0717-4d61-b5d9-2e3f16e79edc
```

- 7 (Optional) Confirm the import task completed successfully.

If the image is large and requires a lot of time, you can exit the utility safely without affecting the operation and check the task status later.

Note You must know the task ID to be able to check the status.

```
glance --os-image-api-version 2 task-show <task_id>
```

For example:

```
glance --os-image-api-version 2 task-show 5cdc4a04-5c68-4b91-ac44-37da07ec82ec
+-----+-----+
| Property | Value |
+-----+-----+
| created_at | 2015-10-15T21:20:59Z |
| expires_at | 2015-10-17T21:21:14Z |
| id | 5cdc4a04-5c68-4b91-ac44-37da07ec82ec |
| input | {"image_properties": {"container_format": "bare", "name": "cirros-  
img"}, "import_from_format": "qcow2", "import_from": "https://launchpad.net/  
cirros/trunk/0.3.0/+download/cirros-0.3.0-x86_64-disk.img"} |
| message | |
| owner | def459fd05d7490e9fda07dbe6ee2d76 |
| result | {"image_id": "2120de75-0717-4d61-b5d9-2e3f16e79edc"} |
| status | success |
| type | import |
| updated_at | 2015-10-15T21:21:14Z |
+-----+-----+
```

- 8 (Optional) Confirm that the import process was successful.

You must know the image ID created by the glance-import command to confirm the import.

```
glance image-show <image_id>
```

The command returns details about the specified image.

- 9 (Optional) Confirm the image is included in the Image Service.

```
glance image-list
```

The command returns a list of all images that are available in the Image Service.

Modify Image Settings

After an image is loaded, you can modify the image settings, such as image name, description, and the public and protected settings.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Images**.
- 4 Select the image to edit.
- 5 In the Actions column, click **Edit Images**.
- 6 Modify the settings as necessary.
- 7 Click **Update Image**.

The Images page redisplay with the changed information.

Modify Image Resource Metadata

After an image is loaded, you can modify the image resource metadata settings by adding or removing metadata tags for the image definition. Image resource metadata can help end users determine the nature of an image, and is used by associated OpenStack components and drivers that interface with the Image Service.

You can manage metadata definitions on the Metadata Definitions page located at **Admin > System > Metadata Definitions**.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Images**.
- 4 Click the image to modify.
- 5 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.

6 Modify the settings as necessary.

The Update Metadata dialog has two columns. The right shows metadata tags already applied to the image, and the left column shows available metadata tags, which are grouped by category, such as Guest Customization, Instance Config Data, and so on.

a To add a metadata tag to the image definition, click the plus sign (+).

The item moves to the Existing Metadata column and is highlighted.

b Enter the metadata value in the provided field, if applicable.

c To remove a metadata tag from the image definition, click the minus sign (-).

7 Click **Save**.

Configuring Images for Windows Guest Customization

You can configure images for Windows guest customization directly in the VMware Integrated OpenStack dashboard by applying the guest customization metadata to the Glance image used to create an instance.

The Windows guest customization feature provides an alternative to the cloudbase-init approach to enabling guest customization. If an image currently uses cloudbase-init, do not use the VMware Integrated OpenStack Windows guest customization feature.

Prerequisites

- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- Verify that you have an appropriate Windows OS image available in the Glance Image Service.
- Verify that the correct versions of the Microsoft System Preparation tools (sysprep) for each guest operating system you want to customize are installed in vSphere. See [Installing the Microsoft Sysprep Tool](#) in the vSphere product documentation.
- Verify that VMware Tools is installed on the source image.
- Verify that the image disk type property correctly reflects the image disk type prior to import.

This applies only to images imported into Glance in VMware Integrated OpenStack versions earlier than 2.0. In version 2.0.x and later, image properties (such as disk type) are automatically introspected during the Glance import process.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.

- 3 (Optional) Preview the Guest Customization Options metadata definition.
 - a Select **Admin > System > Metadata Definitions**.
 - b Click **Guest Customization Options**.
 - c Click the **Contents** tab.

You can only view the metadata definitions in the VMware Integrated OpenStack dashboard. You cannot modify the metadata.

- 4 Select **Admin > System > Images**.
- 5 Locate the Windows image to modify.
- 6 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
- 7 In the column under Available Metadata, expand the **Guest Customization Options** tab.

Note If the **Guest Customization Options** tab is not present, the related metadata properties might already be configured.

- 8 Click the plus sign (+) next to the guest customization option you want to add.



Tip You can add all the options simultaneously by clicking the plus sign (+) on the top **Guest Customization Options** tab.

In the column under Existing Metadata, the newly added metadata properties appear.

Note You may need to scroll to the bottom of this column to see the newly added metadata properties.

- 9 Configure the metadata properties.

Metadata Property	Description
Auto logon count	Applies the windows_logon_count metadata property. Enter the number of times the machine can automatically logged in to as Administrator . Typically, this value is set to 1, but you can increase the value if your configuration requires multiple reboots. This value might be determined by the list of commands executed by the GuiRunOnce command.
Automatic logon	Applies the windows_auto_logon metadata property. If selected, the VM is automatically logged in to as Administrator.
Maximum number of connections	Applies the windows_max_connect metadata property. Enter the number of client licenses purchased for the Windows server being installed. Note This property is applied only if the windows_license_mode metadata property, described below, is set to PerServer.

Metadata Property	Description
Product Key	<p>Applies the windows_product_key metadata property.</p> <p>Enter a valid serial number which is included in the answer file when mini-setup runs.</p> <p>Note This serial number is ignored if the original guest operating system was installed using a volume-licensed CD.</p>
Server licensing mode	<p>Applies the windows_license_mode metadata property.</p> <p>Select the licensing mode that matches your source image: PerServer or PerSeat.</p>
Windows workgroup to join	<p>Applies the windows_join_workgroup metadata property.</p> <p>Select the workgroup that the VM should join.</p>

10 Click **Save**.

The image metadata is now configured for Windows guest customization and are applied for all future VMs that are created from this image.

Configure QoS Resource Allocation for Instances Using Image Metadata

You can control the QoS resource allocations, such as limits, reservations, and shares, for CPU, RAM, disk IOPS, and virtual network interface (VIF) by modifying the metadata of the source image used to create the instance. All instances subsequently created from the image inherit the metadata settings.

QoS resource allocation for an instance can also be specified by flavor metadata. In the event of a conflict, the image metadata configuration overrules the flavor metadata configuration. See [Configure QoS Resource Allocation for Instances Using Flavor Metadata](#).

Prerequisites

- Requires VMware Integrated OpenStack version 2.0.x or later.
- Requires vSphere version 6.0 or later.
- Verify that VMware Integrated OpenStack is running in vSphere.
- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Images**.
- 4 Click the image to modify.
- 5 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.

- 6 In the column under Available Metadata, expand the **VMware Quota** tab.

Note If the **VMware Quota** tab is not present, the related metadata properties might already be configured.

- 7 Click the plus sign (+) next to the VMware Quota metadata property you want to add.



Tip You can add all the options simultaneously by clicking the plus sign (+) on the **VMware Quota** tab.

In the column under Existing Metadata, the newly added metadata properties appear .

- 8 Configure the metadata properties.

Metadata Property	Description
Quota: CPU Limit	Applies the <code>quota_cpu_limit</code> metadata property. Specifies the upper limit for CPU allocation in MHz. This parameter ensures that the instance never uses more than the defined amount of CPU allocation. Enter 0 for unlimited CPU allocation.
Quota: CPU Reservation	Applies the <code>quota_cpu_reservation</code> metadata property. Specifies the guaranteed minimum CPU reservation in MHz. This parameter ensures that the instance has the reserved amount of CPU cycles available during resource contention.
Quota: CPU Shares Level	Applies the <code>quota_cpu_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota_cpu_shares_value</code> metadata property. See Quota: CPU Shares Value below.
Quota: CPU Shares Value	Applies the <code>quota_cpu_shares_value</code> metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the <code>quota_cpu_shares_level</code> metadata property to custom . Otherwise this property is ignored.
Quota: Disk IO Limit	Applies the <code>quota_disk_io_limit</code> metadata property. Specifies the upper limit for disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance never uses more than the defined amount of disk IOPS, and can be used to enforce a limit on the instance's disk performance. Enter 0 for unlimited IOPS.
Quota: Disk IO Reservation	Applies the <code>quota_disk_io_reservation</code> metadata property. Specifies the guaranteed minimum disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance receives the reserved amount of disk IOPS during resource contention.
Quota: Disk IO Shares Level	Applies the <code>quota_disk_io_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota_disk_io_shares_value</code> metadata property (Quota: Disk IO Shares Value).

Metadata Property	Description
Quota: Disk IO Shares Value	Applies the <code>quota_disk_io_shares_share</code> metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the <code>quota_disk_io_shares_level</code> metadata property to custom . Otherwise this property is ignored.
Quota: Memory Limit	Applies the <code>quota_memory_limit</code> metadata property. Specifies the upper limit for memory allocation in MB. This parameter ensures that the instance never uses more than the defined amount of memory. Enter 0 for unlimited memory allocation.
Quota: Memory Reservation	Applies the <code>quota_memory_reservation</code> metadata property. Specifies the guaranteed minimum memory reservation in MB. This parameter ensures that the instance receives the reserved amount of memory during resource contention.
Quota: Memory Shares Level	Applies the <code>quota_memory_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota_memory_shares_share</code> metadata property (Quota: Memory Shares Value).
Quota: Memory Shares Value	Applies the <code>quota_memory_shares_share</code> metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the <code>quota_memory_shares_level</code> metadata property to custom . Otherwise this property is ignored.
Quota: VIF Limit	Applies the <code>quota_vif_limit</code> metadata property. Specifies the upper limit for VIF bandwidth in Mbps. This parameter ensures that the VIF never uses more than the defined amount of bandwidth. Enter 0 for unlimited bandwidth allocation.
Quota: VIF Reservation	Applies the <code>quota_vif_reservation</code> metadata property. Specifies the guaranteed minimum bandwidth for VIF in Mbps. This parameter ensures that the virtual adapter on the instance gets the reserved amount of bandwidth during resource contention. If the instance uses less than the reserved amount, the remainder is available to other virtual adapters.
Quota: VIF Shares Level	Applies the <code>quota_vif_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota_vif_shares_share</code> metadata property (Quota: VIF Shares Value).
Quota: VIF Shares Value	Applies the <code>quota_vif_shares_share</code> metadata property. in the event that 'custom' is used, this is the number of shares.

9 Click **Save**.

The image metadata is now configured for limits, reservations, and shares for CPU, IOPS, memory, and network bandwidth. This configuration is applied to all future OpenStack instances that are created from this image.

Delete an Existing Image

Deleting an image is permanent and cannot be reversed. You must have administrative permissions to delete an image.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System Panel > Images**.
- 4 Select one or more images to delete.
- 5 Click **Delete Images**.
- 6 Confirm the deletion at the prompt.

Migrating Images

You can migrate images between datastores in a way that preserves their UUID and metadata.

This process requires you to copy the images folder from the current datastore to an offline datastore and modify the Image Service image location configuration.

Procedure

- 1 [Copy the Images Folder to the New Datastore](#)
Copy the images folder, including folder name and relative path, to the new datastore.
- 2 [Update the Migrated Images Location Data](#)
After you copy the images folder to a new datastore, you must update the locations setting on each image to reflect the new datastore.

Copy the Images Folder to the New Datastore

Copy the images folder, including folder name and relative path, to the new datastore.

Prerequisites

Verify that both the current and destination datastores are available.

Procedure

- 1 Using SSH, log in to the ESXi host on which the current Image Service datastore is mounted.
- 2 Switch to root user.

```
sudo su -
```

- 3 Locate the images folder.

The images folder is typically called `images` and resides on the top level.

- Using the `cp` or `scp` Linux command, copy the images folder to the new datastore.

Important When copying the folder to the new datastore, retain both the images folder name and relative path.

What to do next

You must now modify the images data to reflect the new location. For details, see [Update the Migrated Images Location Data](#).

Update the Migrated Images Location Data

After you copy the images folder to a new datastore, you must update the locations setting on each image to reflect the new datastore.

Prerequisites

- Verify that the images folder has been copied to the new datastore.
- Verify that the image folder name and relative path on the new datastore are the same as as on the previous datastore.
- Verify that you know the image ID values of the images you want to update.

Procedure

- Repeat this procedure for all images that you want to migrate.
- Using SSH, log in as administrator to the VMware Integrated OpenStack manager.
- Using SSH, log in to the controller01 node.
- Switch to root user.

```
sudo su -
```

- Execute the `cloudadmin.rc` file.

```
source cloudadmin.rc
```

- (Optional) View a list of images.

```
glance image-list
```

- (Optional) Get the location of a specific image.

Note You must know the image ID to specify the image.

```
glance --os-image-api-version 2 image-show <image_id>
```

The image location is the URL indicated by the `locations` parameter.

```
vsphere://<vcenter_ip>/folder/<image_folder_name>/<image_id>dcPath=<path_to_datacenter>&dsName=<old_datastore_name>
```

For example:

```
vsphere://10.20.123.456/folder/images/6c4a7e0d-65e7-4f3c-9dde-0de75f729a0c
?dcPath=Datacenter1&dsName=old_ds
```

- 8 Update the image's location URL to reflect the destination datastore to complete the migration of a single image.

- a Add the new location to the image configuration.

```
glance --os-image-api-version 2 location-add <image_id> --url <new_url>
```

Option	Description
image_id	Specifies the image to be modified.
new_url	The new URL is the same as the previous URL except the dsName argument specifies the name of the new datastore. vsphere://<vcenter_ip>/folder/<image_folder_name>/<image_id>dcPath=<path_to_datacenter>&dsName=<new_datastore_name>

If the command returns a 400 Bad Request: Invalid Location message, verify that file path of the image on the destination datastore is correct.

- b Remove the old location from the image configuration.

```
glance --os-image-api-version 2 location-delete <image_id> --url <old_url>
```

- c View the image information again to verify that the locations parameter correctly reflects the new datastore.

```
glance --os-image-api-version 2 image-show <image_id>
```

The image is successfully migrated.

Add a VM Template as an Image

You can add existing VM templates to your VMware Integrated OpenStack deployment as Glance images. This enables users to boot instances, create bootable block storage volumes, and other functions available to Glance images.

Prerequisites

- Verify that the existing VMs template resides in the same vCenter as your VMware Integrated OpenStack deployment.
- Verify that the following conditions do apply.
 - The VM template does not have multiple disks.

- The VM template does not have a CD-ROM drive.
- The VM template does not have a floppy disk drive.

Procedure

1 Prepare the VM template.

Configure the metadata settings as necessary.

- The `vmware_ostype` is required for Windows images, but optional for Linux images.
- The `hw_vif_model` is recommended for specifying NIC type. Before defining this setting, confirm the correct NIC type for this image template. For example, if this setting is undefined, the instance is provisioned with the E1000 NIC by default. To ensure another NIC is provisioned, define this setting appropriately.

For example, to provision the VMXNET3 NIC, the metadata definition is

`hw_vif_model=VirtualVmxnet3`.

- The following metadata settings are not required.
 - `vmware_adaptype`
 - `vmware_disktype`

2 Log in to the OpenStack management cluster.

3 Run the `glance` command to obtain, define, and import the image.

```
glance image-create --name <NAME> \
  --disk-format vmdk --container-format bare
  --property vmware_ostype=ubuntu64Guest
  --property hw_vif_model=VirtualVmxnet3

glance location-add <glance_image_UUID> --url "vi://<vcenter-host>/<datacenter-path>/vm/<sub-
folders>/<template_name> IMAGE_ID"
```

The `location-add` command points to the inventory path for the VM template and can refer to either VM or host. For example:

```
"vi://<datacenter-path>/vm/<template_name>"
or
"vi://<datacenter-path>/host/<host_name>/<template_name>"
```

The `vm` and `host` keywords in the inventory path represent the **VM and Templates View** and **Host and Cluster View** hierarchy in your vSphere Web Client.

Configure Images to Enable Live Resize of VMs Deployed From That Image

Starting with VMware Integrated OpenStack 4.0, you can deploy VMs from images that can later be resized with no need to power them off. You can change memory, vCPU, and root disk sizes.

The live-resize functionality uses the `os_live_resize` property for images that is not available in previous versions of VMware Integrated OpenStack, so that you must add it to your existing images to be able to resize new VMs without powering them off. The value of `os_live_resize` can be `memory`, `disk`, and `vcpu`, or any combination separated by commas. For example `os_live_resize=disk,memory,vcpu`.

Prerequisites

To be able to deploy VMs that are capable of live-resizing, the following requirements for the image must be met.

- Create the VM image in VMware Integrated OpenStack 4.0 or later, so that the `os_live_resize` property is available.
- To be able to resize disks, deploy the VMs as full clones not linked clones and use SCSI virtual disk adapter types. IDE adapter type is not supported.

Procedure

1 Log in to the OpenStack management cluster.

2 Create new image that uses a Virtual Machine Disk.

```
openstack image create --container-format bare --disk-format vmdk \
```

3 Configure a SCSI virtual disk for the `vmware_adaptertype` property.

```
--property vmware_adaptertype="lsiLogicsas" --property vmware_disktype="sparse" \
```

4 Configure the VMs to deploy from the image to be full clones and not linked ones.

```
--property vmware_ostype="otherGuest64" --property img_linked_clone=False \
```

5 Configure the VM settings to be available for live-resize through the `os_live_resize` property.

```
--property os_live_resize=vcpu,memory,disk --file cirros-0.3.5-x86_64-disk.vmdk  
--public cirros
```

You created a new image for VMs that can be resized with no need to power them off.

Modify the Default Behavior for Nova Snapshots

By default, Nova snapshots are Glance images that are stored and organized as VM templates in the vCenter configured for VMware Integrated OpenStack. You can modify this behavior so that snapshots are stored as streamOptimized VMDK disks instead.

Before VMware Integrated OpenStack 2.5, the default behavior was to store Nova snapshots as streamOptimized VMDK disks. This procedure enables you to restore the pre-2.5 default.

Procedure

- 1 Implement the custom.yml file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample /opt/vmware/vio/custom/custom.yml
```

- 2 Open the /opt/vmware/vio/custom/custom.yml file in a text editor.
 - a Uncomment the nova_snapshot_format parameter.
 - b Change the setting to **streamOptimized**.

```
#####
# Glance Template Store
# options that affect the use of glance template store
#####
#glance_default_store: vi
nova_snapshot_format: streamOptimized
#cinder_image_format: template
```

- 3 Save the custom.yml file.
- 4 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment configure
```

Note Pushing the configuration briefly interrupts OpenStack services.

Modify the Default Cinder Upload-to-Image Behavior

By default, the Block Storage upload-to-image feature creates a Glance image from a Cinder volume that is stored and organized as a VM template. You can modify this behavior so that the images are stored as streamOptimized VMDK disks instead.

Before VMware Integrated OpenStack 4.0, the default behavior was to store the Glance images as streamOptimized VMDK disks. This procedure enables you to restore the pre-4.0 default.

Procedure

- 1 Implement the custom.yml file.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample
/opt/vmware/vio/custom/custom.yml
```

- 2 Open the /opt/vmware/vio/custom/custom.yml file in a text editor.
 - a Uncomment the cinder_image_format parameter.
 - b Change the setting to **streamOptimized**.

```
#####
# Glance Template Store
# options that affect the use of glance template store
#####
#glance_default_store: vi
#nova_snapshot_format: template
cinder_image_format: streamOptimized
```

- 3 Save the custom.yml file.
- 4 Push the new configuration to your VMware Integrated OpenStack deployment.

```
viocli deployment configure
```

Note Pushing the configuration briefly interrupts OpenStack services.

Working with Flavors

In OpenStack, a flavor is a preset configuration that defines the compute, memory, and storage capacity of an instance. When you create an instance, you configure the server by selecting a flavor. Administrative users can create, edit, and delete flavors.

Do not delete any of the default flavors.

This section includes the following topics:

- [Default Flavor Configurations](#)
- [Create a Flavor](#)
- [Delete a Flavor](#)
- [Modify Flavor Metadata](#)
- [Configure QoS Resource Allocation for Instances Using Flavor Metadata](#)

Default Flavor Configurations

The default OpenStack deployment provides five default flavors ranging from tiny to extra large.

Name	vCPUs	RAM (MB)	Disk (GB)
m1.tiny	1	512	1
m1.small	1	2048	20
m1.medium	2	4096	40
m1.large	4	8192	80
m1.xlarge	8	16384	160

Create a Flavor

Administrative users can create custom flavors.

Prerequisites

Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 On the VMware Integrated OpenStack dashboard, select the admin project from the drop-down menu in the title bar.
- 2 Select **Admin > System Panel > Flavors**.
- 3 Click **Create Flavor**.
- 4 In the Create Flavor dialog box, configure the new flavor.

Parameter	Description
Name	Name for the flavor.
ID	Integer or a UUID4 value that identifies the flavor. If this parameter is left blank or has a value of auto , OpenStack automatically generates a UUID.
VCPUs	Number of virtual CPUs that an instance made from this flavor will use.
RAM MB	Megabytes of RAM for virtual machines made from this flavor.
Root Disk GB	Gigabytes of disk used for the root (/) partition in instances made from this flavor.
Ephemeral Disk GB	Gigabytes of disk space to use for the ephemeral partition. If unspecified, the value is 0 by default. Ephemeral disks offer machine local disk storage linked to the life cycle of a VM instance. When a VM is terminated, all data on the ephemeral disk is lost. Ephemeral disks are not included in snapshots.
Swap Disk MB	Megabytes of swap space to use. If unspecified, the default is 0.

- 5 Click **Create Flavor** at the bottom of the dialog box to complete the process.
- 6 (Optional) Specify which projects can access instances created from specific flavors.
 - a On the Flavors page, click **Edit Flavor** in the Actions column of the instance.
 - b In the Edit Flavor dialog box, click the **Flavor Access** tab.
 - c Use the toggle controls to select the projects that can access the instance.
 - d Click **Save**.
- 7 (Optional) Modify the settings of a specific flavor.
 - a On the Flavors page, click **Edit Flavor** in the Actions column of the instance.
 - b In the Edit Flavor dialog box, modify the settings in either the **Flavor Info** or **Flavor Access** tab.
 - c Click **Save**.

Delete a Flavor

You can manage the number and variety of flavors by deleting those that no longer meet users' needs, duplicate other flavors, or for other reasons.

Note You cannot undo the deletion of a flavor. Do not delete default flavors.

Prerequisites

You must be logged in to the VMware Integrated OpenStack dashboard as a cloud administrator to perform this task.

Procedure

- 1 In the VMware Integrated OpenStack dashboard, select the admin project from the drop-down menu in the title bar.
- 2 Select **Admin > System Panel > Flavors**.
- 3 Select the flavors to delete.
- 4 Click **Delete Flavors**.
- 5 At the prompt, confirm the deletion.

Modify Flavor Metadata

You can modify the metadata of a flavor to dynamically add properties to all the instances that are subsequently created that use that flavor.

You can also use image metadata to specify many flavor metadata settings. If a conflict occurs, the image metadata configuration overrules the flavor metadata configuration.

Prerequisites

- Requires VMware Integrated OpenStack version 2.0.x or greater.
- Requires vSphere version 6.0 or greater.
- Verify that VMware Integrated OpenStack is running in vSphere.
- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Flavors**.
- 4 (Optional) Create a flavor specific to the intended use of the metadata application.
Create a custom flavor to contain the specific configuration. The custom flavor leaves the original flavor configuration intact and available for other instance creation.
- 5 Select the flavor to modify.
- 6 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
- 7 Click the plus sign (+) next to the metadata properties to add.

In the column under Existing Metadata, the newly added metadata properties appear.

8 Configure the metadata properties.

For example, you might have to select an option from a drop-down list or enter a string value.

9 Click **Save**.

The newly added flavor metadata properties are now configured. This configuration is applied to all future OpenStack instances that are created from this flavor.

Configure QoS Resource Allocation for Instances Using Flavor Metadata

You can control the QoS resource allocations, such as limits, reservations, and shares, for CPU, RAM, disk IOPS, and virtual network interface (VIF) by modifying the metadata of the flavor used to create the instance. All instances subsequently created using the flavor inherit the metadata settings.

QoS resource allocation can also be specified by image metadata. In the event of a conflict, the image metadata configuration overrules the flavor metadata configuration. See [Configure QoS Resource Allocation for Instances Using Image Metadata](#).

Prerequisites

- Requires VMware Integrated OpenStack version 2.0.x or greater.
- Requires vSphere version 6.0 or greater.
- Verify that VMware Integrated OpenStack is running in vSphere.
- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Flavors**.
- 4 (Optional) Create a flavor specific to the set of QoS resource allocations.

You must create a custom flavor to contain the specific configuration. This leaves the original flavor configuration intact and available for other uses.

- 5 Select the flavor to modify.
- 6 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
- 7 In the column under Available Metadata, expand the **VMware Quota** tab.

Note If the VMware Quota tab is not present, the related metadata properties might already be configured.

- 8 Click the plus sign (+) next to the VMware Quota metadata property you want to add.



Tip You can add all the options simultaneously by clicking the plus sign (+) on the VMware Quota tab.

In the column under Existing Metadata, the newly added metadata properties appear.

- 9 Configure the metadata properties.

Metadata Property	Description
Quota: CPU Limit	Applies the <code>quota:cpu_limit</code> metadata property. Specifies the upper limit for CPU allocation in MHz. This parameter ensures that the instance never uses more than the defined amount of CPU allocation. Enter 0 for unlimited CPU allocation.
Quota: CPU Reservation	Applies the <code>quota:cpu_reservation</code> metadata property. Specifies the guaranteed minimum CPU reservation in MHz. This parameter ensures that the instance has the reserved amount of CPU cycles available during resource contention.
Quota: CPU Shares Level	Applies the <code>quota:cpu_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota:cpu_shares_value</code> metadata property. See Quota: CPU Shares Value below.
Quota: CPU Shares Value	Applies the <code>quota:cpu_shares_value</code> metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the <code>quota:cpu_shares_level</code> metadata property to custom . Otherwise this property is ignored.
Quota: Disk IO Limit	Applies the <code>quota:disk_io_limit</code> metadata property. Specifies the upper limit for disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance never uses more than the defined amount of disk IOPS, and can be used to enforce a limit on the instance's disk performance. Enter 0 for unlimited IOPS.
Quota: Disk IO Reservation	Applies the <code>quota:disk_io_reservation</code> metadata property. Specifies the guaranteed minimum disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance receives the reserved amount of disk IOPS during resource contention.
Quota: Disk IO Shares Level	Applies the <code>quota:disk_io_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota:disk_io_shares_share</code> metadata property (Quota: Disk IO Shares Value).
Quota: Disk IO Shares Value	Applies the <code>quota:disk_io_shares_share</code> metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the <code>quota:disk_io_shares_level</code> metadata property to custom . Otherwise this property is ignored.

Metadata Property	Description
Quota: Memory Limit	Applies the <code>quota:memory_limit</code> metadata property. Specifies the upper limit for memory allocation in MB. This parameter ensures that the instance never uses more than the defined amount of memory. Enter 0 for unlimited memory allocation.
Quota: Memory Reservation	Applies the <code>quota:memory_reservation</code> metadata property. Specifies the guaranteed minimum memory reservation in MB. This parameter ensures that the instance receives the reserved amount of memory during resource contention.
Quota: Memory Shares Level	Applies the <code>quota:memory_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota:memory_shares_share</code> metadata property (Quota: Memory Shares Value).
Quota: Memory Shares Value	Applies the <code>quota:memory_shares_share</code> metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the <code>quota:memory_shares_level</code> metadata property to custom . Otherwise this property is ignored.
Quota: VIF Limit	Applies the <code>quota:vif_limit</code> metadata property. Specifies the upper limit for VIF bandwidth in Mbps. This parameter ensures that the VIF never uses more than the defined amount of bandwidth. Enter 0 for unlimited bandwidth allocation.
Quota: VIF Reservation	Applies the <code>quota:vif_reservation</code> metadata property. Specifies the guaranteed minimum bandwidth for VIF in Mbps. This parameter ensures that the virtual adapter on the instance gets the reserved amount of bandwidth during resource contention. If the instance uses less than the reserved amount, the remainder is available to other virtual adapters.
Quota: VIF Shares Level	Applies the <code>quota:vif_shares_level</code> metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the custom level is selected, you must include the <code>quota:vif_shares_share</code> metadata property (Quota: VIF Shares Value).
Quota: VIF Shares Value	Applies the <code>quota:vif_shares_share</code> metadata property. in the event that 'custom' is used, this is the number of shares.

10 Click **Save**.

The flavor metadata is now configured for limits, reservations, and shares for CPU, IOPS, memory, and network bandwidth. This configuration is applied to all future OpenStack instances that are created from this flavor.

Working with VMware Integrated OpenStack Carrier Edition



VMware Integrated OpenStack Carrier Edition is an OpenStack distribution that accelerates the deployment of production Network File Virtualization (NFV) services on OpenStack.

As a communication service provider (CSP), you can use VMware Integrated OpenStack Carrier Edition to run Carrier Grade OpenStack solutions on the vCloud NFV platform.

This section includes the following topics:

- [Use Tenant Virtual Data Center for Secure Multi-Tenancy and Resource Allocation](#)
- [Configure Passthrough Devices for Instances](#)

Use Tenant Virtual Data Center for Secure Multi-Tenancy and Resource Allocation

Using the Tenant Virtual Data Center available with the VMware Integrated OpenStack Carrier Edition, you can create virtual datacenters for tenants under different compute nodes that offer specific service level agreements for each telecommunication workload.

While quotas on projects set limits on the OpenStack resources across multiple compute nodes or availability zones, they do not guarantee resource reservation. By using the Tenant Virtual Data Center to allocate CPU and memory for an OpenStack project or tenant on a compute node, you provide a resource guarantee for tenants and avoid noisy neighbor scenarios in a multi-tenant environment.

To manage the Tenant Virtual Data Center, you use the VMware Integrated OpenStack command line utility. The `viocli inventory-admin` command includes parameters to create, list, and delete a tenant virtual data center.

To use the Tenant Virtual Data Center to allocate resources at the compute node level. To allocate resources for Virtual Network Functions (VNF), see [Configure QoS Resource Allocation for Instances Using Flavor Metadata](#).

Prerequisites

Verify that VMware Integrated OpenStack version 4.0 or later is deployed and running.

Procedure

- 1 Using SSH, log in to the VMware Integrated OpenStack manager.

2 Create the Tenant Virtual Data Center.

```
viocli inventory-admin create-tenant-vdc
  --project-id <project-id>
  --compute <compute-node>
  --name <vdc-name>
  --cpu-limit <cpu-limit>
  --cpu-reserve <cpu-reserve>
  --mem-limit <mem-limit>
  --mem-reserve <mem-reserve>
```

Parameter	Description
project-id	OpenStack project ID
compute-node	Compute node on VMware Integrated OpenStack
vdc-name	Name of the Tenant Virtual Data Center
cpu-limit	Upper limit for CPU in MHz within the compute node
cpu-reserve	Minimum guaranteed or reserved value for CPU in MHz within the compute node
mem-limit	Upper limit for memory in MB within the compute node
mem-reserve	Minimum guaranteed or reserved value for memory in MB within the compute node

The following example includes typical values.

```
viocli inventory-admin create-tenant-vdc
  --project-id 908909ca3db4460faaa0f765757470ac
  --compute compute01
  --name computeA_gold
  --cpu-limit 10000
  --cpu-reserve 8000
  --mem-limit 20000
  --mem-reserve 10000
```

3 Get the UUID for the Tenant Virtual Data Center.

```
viocli inventory-admin list-tenant-vds
```

The result lists the Tenant Virtual Data Center name and its UUID.

```
+-----+-----+
| name                                     | id                                     |
+-----+-----+
| computeA_gold (4c238c45dbcb433fb6105420c3b05b63) | 4c238c45dbcb433fb6105420c3b05b63 |
+-----+-----+
```

4 Create the flavor.

```
openstack flavor create
  --disk <disk-size-gb>
  --ram <memory-in-mb>
  --vcpus <vcpu-count>
  --private
  --project <project>
  <flavor-name>
```

Parameter	Description
disk-size-gb	Disk size in GB (default 0G)
memory-in-mb	Memory size in MB (default 256M)
vcpu-count	Number of vcpus (default 1)
project	Project allowed to access private flavor specified as name or ID. Must be used with --private option.
flavor-name	New flavor name

The following example includes typical values.

```
openstack flavor create
  --disk 10
  --ram 2048
  --vcpus 1
  --private
  --project 908909ca3db4460faaa0f765757470ac
  companyA_gold
```

The result lists the flavor name and its UUID.

```
+-----+-----+
| Field                | Value                |
+-----+-----+
| OS-FLV-DISABLED:disabled | False                |
| OS-FLV-EXT-DATA:ephemeral | 0                    |
| disk                  | 10                   |
| id                    | 7796b6ea-17b1-4dec-868c-12e4a7196efd |
| name                  | companyA_gold       |
| os-flavor-access:is_public | False                |
| properties            |                      |
| ram                   | 2048                 |
| rxtx_factor           | 1.0                  |
| swap                  |                      |
| vcpus                 | 1                    |
+-----+-----+
```

5 Use the OpenStack flavor extra specs to expose the Tenant Virtual Data Center.

```
openstack flavor set
  --property vmware:tenant_vdc=<UUID-TvDC> <UUID-flavor>
```

The following example includes values for the Tenant Virtual Data Center and flavor UUIDs.

```
openstack flavor set
  --property vmware:tenant_vdc=4c238c45dbcb433fb6105420c3b05b63
  7796b6ea-17b1-4dec-868c-12e4a7196efd
```

6 (optional) You can use the following command to expose the VNF level resource allocation with the Tenant Virtual Data Center.

```
openstack flavor set
  --property vmware:tenant_vdc=<UUID-TvDC>
  --property quota:memory_reservation_percent <memory-percent>
  --property quota:cpu_reservation_percent <cpu-percent> <UUID-flavor>
```

The following example includes typical values.

```
openstack flavor set
  --property vmware:tenant_vdc=4c238c45dbcb433fb6105420c3b05b63
  --property quota:memory_reservation_percent 100
  --property quota:cpu_reservation_percent 100 7796b6ea-17b1-4dec-868c-12e4a7196efd
```

What to do next

When the Tenant Virtual Data Center is no longer needed, delete it.

```
viocli inventory-admin delete-tenant-vdc --id <UUID-TvDC>
```

Configure Passthrough Devices for Instances

You can create instances that use the Single Root I/O Virtualization (SR-IOV) specification by modifying the metadata parameters of the flavor and image used to create the instance. SR-IOV is a specification that allows a single Peripheral Component Interconnect Express (PCIe) physical device under a single root port to appear as multiple separate physical devices.

For more information about SR-IOV requirements and supported features, see the vSphere Web Client documentation.

The following table describes the key components of SR-IOV and their role.

Table 8-1. SR-IOV Components in the VMware Integrated OpenStack context

Component	Role
Nova Compute	<ul style="list-style-type: none"> Collects the list of SR-IOV devices and updates the list of PCI device specifications. Embeds the host object ID in device specifications.
Nova PCI manager	<ul style="list-style-type: none"> Creates and maintains a device pool with address, vendor ID, product ID, and host ID. Allocates and deallocates PCI devices to instances based on PCI requests.
Nova scheduler	<ul style="list-style-type: none"> Schedules instance placement on hosts that matches the PCI requests
vSphere	<ul style="list-style-type: none"> Manages hosts in a dedicated compute cluster with NICs and hosts enabled for SR-IOV. <p>A separate compute cluster is recommended because DRS rules do not work on devices enabled for SR-IOV.</p>

Prerequisites

- Verify that your deployment is VDS based. SR-IOV does not work with NSX.
- Requires VMware Integrated OpenStack version 2.0.x or greater.
- Requires vSphere version 6.0 or greater.

Procedure**1** [Enabling SR-IOV on Network Adapters in vSphere](#)**2** [Configure GPU Passthrough Devices for OpenStack Instances](#)

Starting with VMware Integrated OpenStack 3.1, you can create OpenStack instances that use GPU physical functions (enabled using directpath I/O) or virtual functions (SR-IOV) from vSphere.

3 [Configure Network DirectPath I/O Passthrough for OpenStack Instances](#)

Starting with VMware Integrated OpenStack 3.1, you can create OpenStack instances that use network physical functions with the DirectPath I/O technology from VMware.

4 [Modify Flavor Metadata to Enable SR-IOV](#)

You must modify the metadata for a flavor to enable SR-IOV. All instances created from a SR-IOV-enabled flavor and SR-IOV-enabled image inherit the SR-IOV property.

5 [Modify Image Metadata to Enable SR-IOV](#)

You must modify the metadata for an image to enable SR-IOV. All instances created from a SR-IOV-enabled flavor and SR-IOV-enabled image inherit the SR-IOV property.

Enabling SR-IOV on Network Adapters in vSphere

Procedure

- 1 In the vSphere Web Client, navigate to the host.
- 2 On the **Configure** tab, expand **Networking** and select **Physical adapters**.
You can look at the SR-IOV property to see whether a physical adapter supports SR-IOV.
- 3 Select the physical adapter and click **Edit adapter settings**.
- 4 Under SR-IOV, select **Enabled** from the **Status** drop-down menu.
- 5 In the **Number of virtual functions** text box, type the number of virtual functions that you want to configure for the adapter.
- 6 Click **OK**.
- 7 Restart the host.

Because DRS rules do not work on SR-IOV-enabled devices, create a dedicated compute cluster for SR-IOV-enabled hosts and adapters, called `vmnics`.

The virtual functions become active on the NIC port represented by the physical adapter entry. They appear in the PCI Devices list in the **Settings** tab for the host.

You can use the `esxcli network sriovnic vCLI` commands to examine the configuration of virtual functions on the host.

What to do next

You can now configure the flavor and image metadata in the VMware Integrated OpenStack dashboard.

Configure GPU Passthrough Devices for OpenStack Instances

Starting with VMware Integrated OpenStack 3.1, you can create OpenStack instances that use GPU physical functions (enabled using directpath I/O) or virtual functions (SR-IOV) from vSphere.

Consumption of GPU and passthrough features is achieved by using the appropriate flavor. Modify the metadata parameters of the flavor to create the instance.

Prerequisites

Make sure that you perform the following settings in your environment before you can configure GPU passthrough devices:

- Enable DirectPath I/O in vSphere. See the *DirectPath I/O* chapter in the *VMware vSphere 6.5 Documentation*.
- Enable SR-IOV on GPU devices on your ESXi hosts. See *Configuring AMD Multiuser GPU Using vDGA* in the VMware Horizon documentation.

Procedure

- 1 Log in to the OpenStack Management Server.
- 2 Create `custom.yml` file if it does not exist.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample
/opt/vmware/vio/custom/custom.yml
```

- 3 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
- 4 Create PCI alias using VIO customization by editing the `custom.yml` file according to your configuration.
 - a Edit `nova_pci_alias` value to create PCI alias based on `device_type`, `vendor_id`, and `product_id` and name the alias, for example:

```
nova_pci_alias: [{"product_id": "692f", "vendor_id": "1002", "device_type:" "type-VF", "name":
"gpu-vf"}]
```

- b Save the `custom.yml` file.
- 5 Push the new configuration to your VMware Integrated OpenStack deployment.

Refresh of the configuration briefly interrupts the OpenStack services.

```
viocli deployment configure --tags nova_api_config
```

What to do next

[Modify Flavor Metadata to Enable SR-IOV.](#)

Configure Network DirectPath I/O Passthrough for OpenStack Instances

Starting with VMware Integrated OpenStack 3.1, you can create OpenStack instances that use network physical functions with the DirectPath I/O technology from VMware.

Consumption of DirectPath I/O passthrough features is achieved by using the appropriate flavor. Modify the metadata parameters of the flavor to create the instance.

Prerequisites

Make sure that you perform the following settings in your environment before you can configure DirectPath I/O passthrough devices:

- Enable DirectPath I/O in vSphere. See the *DirectPath I/O* chapter in the *VMware vSphere 6.5 Documentation*.

Procedure

- 1 Log in to the OpenStack Management Server.

- 2 Create `custom.yml` file if it does not exist.

```
sudo mkdir -p /opt/vmware/vio/custom
sudo cp /var/lib/vio/ansible/custom/custom.yml.sample
/opt/vmware/vio/custom/custom.yml
```

- 3 Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
- 4 Create PCI alias using VIO customization by editing the `custom.yml` file according to your configuration.
 - a Edit `nova_pci_alias` value to create PCI alias based on `device_type`, `vendor_id`, and `product_id` and name the alias, for example:

```
nova_pci_alias: [{"device_type": "type-VF", "name": "sriov"}, {"vendor_id": "15b3",
"product_id": "1013", "device_type": "type-PF", "name": "fpt"}]
```

- b Save the `custom.yml` file.
- 5 Push the new configuration to your VMware Integrated OpenStack deployment.

Refresh of the configuration briefly interrupts the OpenStack services.

```
viocli deployment configure --tags nova_api_config
```

What to do next

[Modify Flavor Metadata to Enable SR-IOV.](#)

Modify Flavor Metadata to Enable SR-IOV

You must modify the metadata for a flavor to enable SR-IOV. All instances created from a SR-IOV-enabled flavor and SR-IOV-enabled image inherit the SR-IOV property.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Flavors**.
- 4 (Optional) Create a flavor dedicated to the SR-IOV specification.
The original flavor configuration remains intact and available for other uses.
- 5 Select the flavor to modify.
- 6 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
- 7 In the column under Available Metadata, expand the **VMware Driver Options for Flavors** tab.

Note If the **VMware Driver Options for Flavors** tab is not present, the related metadata property might already be configured.

- 8 Click the plus sign (+) next to the PCI Passthrough alias metadata property.

In the column under Existing Metadata, the newly added metadata property and the default value appear. The numerical portion represents the number of virtual functions that you can request.

The PCI Passthrough alias refers to a PCI request specification that contains `vendor_id`, `product_id`, and `device_type`. In VMware Integrated OpenStack, the alias is already created and refers to a PCI request specification that you can use to allocate any device regardless of the `vendor_id`, `product_id`, and `device_type`.

- 9 Increase the numerical value as needed.

The maximum number of virtual functions allowed is **10**.

- 10 Click **Save**.

You can now modify image metadata to enable SR-IOV.

Modify Image Metadata to Enable SR-IOV

You must modify the metadata for an image to enable SR-IOV. All instances created from a SR-IOV-enabled flavor and SR-IOV-enabled image inherit the SR-IOV property.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- 2 Select the admin project from the drop-down menu in the title bar.
- 3 Select **Admin > System > Images**.
- 4 (Optional) Create an image definition dedicated to the SR-IOV specification.
The original image configuration remains intact and available for other uses.
- 5 Select the image to modify.
- 6 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
- 7 In the column under Available Metadata, expand the **VMware Driver Options** tab.

Note If the **VMware Driver Options** tab is not present, the related metadata property might already be configured.

- 8 Click the plus sign (+) next to the Virtual Network Interface metadata property.

In the column under Existing Metadata, the newly added metadata property appears as `hw_vif_model`.

- 9 From the drop-down list, select **VirtualSriovEthernetCard**.

- 10 Click **Save**.

VMware Integrated OpenStack CLI Command Reference

9

The VMware Integrated OpenStack CLI command has specific syntax requirements.

This section includes the following topics:

- [viocli backup Command](#)
- [viocli dbverify Command](#)
- [viocli deployment Command](#)
- [viocli ds-migrate-prep Command](#)
- [viocli epops Command](#)
- [viocli inventory-admin Command](#)
- [viocli recover Command](#)
- [viocli restore Command](#)
- [viocli rollback Command](#)
- [viocli services Command](#)
- [viocli show Command](#)
- [viocli upgrade Command](#)
- [viocli volume-migrate Command](#)

viocli backup Command

Use the `viocli backup` command to create a backup of either manager server data or the OpenStack database. This command requires an NFS server to be available for the VMware Integrated OpenStack CL to mount.

The `viocli backup` command uses the following syntax.

```
viocli backup mgmt_server [-d NAME] NFS_VOLUME [-h] [-v]
viocli backup openstack_db [-d NAME] NFS_VOLUME [-h] [-v]
```

Parameter	Mandatory or Optional	Description
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
<i>NFS_VOLUME</i>	Mandatory	Name or IP address of the target NFS volume and directory in the format <i>remote_host:remote_dir</i> . For example: 192.168.1.77:/backups
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

The backup file of the VMware Integrated OpenStack management server is labeled with the time stamp `vio_ms_yyyymmddhhmmss`. The backup file of the VMware Integrated OpenStack database is labeled with the time stamp `vio_os_db_yyyymmddhhmmss`.

viocli dbverify Command

Use the `viocli dbverify` command to check the VMware Integrated OpenStack database for known problems, such as duplicated or missing keys, that can cause problems during the upgrade procedure.

The `viocli dbverify` command uses the following syntax.

```
viocli dbverify [-d NAME] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

viocli deployment Command

Use the `viocli deployment` command to manage your VMware Integrated OpenStack deployment.

The `viocli deployment` command uses the following syntax.

```
viocli deployment ACTION [-d NAME] [-p] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
ACTION	Mandatory	<p>start Start the deployment.</p> <p>stop Stop the deployment.</p> <p>pause Pause the deployment.</p> <p>resume Resume the paused deployment.</p> <p>configure Reconfigure the entire deployment.</p> <p>cert-req-create Create a certificate signing request for a certificate authority.</p> <p>cert-update Update VMware Integrated OpenStack with the provided certificate.</p> <p>getlogs Generate log files for the current deployment, including the ansible executed commands and output. Log files are written to <code>/var/log/viocli/viocli.log</code> and rotated after they reach 100 MB, with a maximum of seven rotations.</p> <p>status Generate reports on the following potential issues:</p> <ul style="list-style-type: none"> ■ Synchronization issues between the management server and OpenStack nodes, including time of occurrence, the affected node, and reason for failed status. ■ Connections to OpenStack processes and average connection count. ■ Broken network connections, including time of occurrence, the hostnames and ports that failed to connect, or single host, if applicable. ■ OpenStack database issues, including time of occurrence, FAILED or SUCCESS status, reason for failure if applicable, and the current size of the database cluster. ■ Missing processes, including time of occurrence, node where issue occurred, status, and reason for failure, if applicable.
Use one of the following positional arguments:		
<ul style="list-style-type: none"> ■ start ■ stop ■ pause ■ resume ■ configure ■ cert-req-create ■ cert-update ■ getlogs ■ status 		
<code>-d, --deployment NAME</code>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
<code>-p, --progress</code>	Optional	Show progress of the current upgrade operation.
<code>-h, --help</code>	Optional	Show the use and arguments for this command.
<code>-v, --verbose</code>	Optional	Enter verbose mode.

viocli ds-migrate-prep Command

Use the `viocli ds-migrate-prep` command to prepare a datastore for maintenance. The `viocli ds-migrate-prep` command helps you ensure that the specified datastore in your VMware Integrated OpenStack deployment does not contain broken references.

The `viocli ds-migrate-prep` command uses the following syntax.

```
viocli ds-migrate-prep [-d NAME] DC_NAME DS_NAME [-h] [-v]
```

Parameter	Mandatory or Optional	Description
<code>-d, --deployment NAME</code>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
<code>DC_NAME</code>	Mandatory	Specifies data center by name.
<code>DS_NAME</code>	Mandatory	Specifies datastore by name.
<code>-h, --help</code>	Optional	Show the use and arguments for this command.
<code>-v, --verbose</code>	Optional	Enter verbose mode.

viocli epops Command

Use the `viocli epops` command to manage the Endpoint Operations Management agent.

VMware Integrated OpenStack has been tested with vRealize Operations Manager 6.2.1

The `viocli Endpoint Operations` command uses the following syntax.

```
viocli epops ACTION [-d NAME] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
ACTION	Mandatory	
Use one of the following positional arguments:		
■ <code>install</code>		install Install the Endpoint Operations Management agent.
■ <code>uninstall</code>		uninstall Uninstall the Endpoint Operations Management agent.
■ <code>config</code>		config Configure the Endpoint Operations Management agent.
■ <code>start</code>		start Start the Endpoint Operations Management agent.
■ <code>stop</code>		stop Stop the Endpoint Operations Management agent.
<code>-d, --deployment NAME</code>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
<code>-h, --help</code>	Optional	Show the use and arguments for this command.
<code>-v, --verbose</code>	Optional	Enter verbose mode.

viocli inventory-admin Command

Use the `viocli inventory-admin` command to compare the compute and block storage inventories against the vSphere inventory and to discover and remove orphaned objects. Orphaned instances are instances that do not have corresponding VMs in OpenStack and vSphere.

The `viocli inventory-admin` command collects vCenter and OpenStack credentials from internal inventories. This command requires that you enter the OpenStack administrative password. To prevent having to enter the password each time, set the `OS_PASSWORD` environment variable.

The `viocli inventory-admin` command uses the following syntax.

```
viocli inventory-admin SHOW_ACTION [-d NAME] [--json] \
  [--pretty] [--all] [--no-grace-period] \
  [--force] [-h] [-v]
```

```
viocli inventory-admin CLEAN_ACTION [-d NAME] [--json] \
  [--pretty] [--all] [--no-grace-period] \
  [--force] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
SHOW_ACTION Use one of the following positional arguments: <ul style="list-style-type: none"> ■ show-instances ■ show-instance-vms ■ show-shadow-vms 	Mandatory	show-instances Display orphaned OpenStack instances. show-instance-vms Display orphaned vSphere instances. show-shadow-vms Display orphaned volume shadow VMs. These are volume VMs that do not have corresponding block storage volumes in the OpenStack database.
CLEAN_ACTION Use one of the following positional arguments: <ul style="list-style-type: none"> ■ clean-instances ■ clean-instance-vms ■ clean-shadow-vms 	Mandatory	clean-instances Remove orphaned OpenStack instances. clean-instance-vms Remove orphaned vSphere instances. clean-shadow-vms Remove orphaned volume shadow VMs. These are volume VMs that do not have corresponding block storage volumes in the OpenStack database.
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
--json	Optional	Return output in JSON format. This is the default format when this command is used noninteractively.
--pretty	Optional	Return output in human-readable format. This is the default format when this command is used interactively.
--all	Optional	Show all objects. The default is to show only orphaned objects.
--no-grace-period	Optional	Disable the default grace period setting. If no grace period is set, the command ignores all objects created or modified in the last 30 minutes.
-f, --force	Optional	Perform operation without requiring confirmation.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

viocli recover Command

Use the `viocli recover` command to recover one node or a group of nodes. Because most OpenStack nodes are stateless, you can recover them without a backup. For OpenStack database nodes, you must have a backup file. An NFS path is required. Use the `viocli show` command to view a detailed list of OpenStack nodes in your deployment.

The `viocli recover` command uses the following syntax.

```
viocli recover [-d [NAME]] <-r ROLE1,ROLE2... | -n NODE1,NODE2...> \
  [-dn BACKUP_NAME] [-nfs NFS_VOLUME] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
<code>-d, --deployment NAME</code>	Automatic	Name of the deployment containing the nodes to be recovered. Applied automatically. The default value is the name of the current deployment.
<code>-r, --role ROLE</code>	Mandatory unless <code>NODE</code> is specified	Recovers all the nodes assigned to a given role. You can specify multiple roles in one command. You can also specify <code>-n, --node</code> to the same command to recover additional nodes that are not assigned to that role. Use the group name as it appears in the VMware Integrated OpenStack manager. To view the group name, select VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name] . The valid role names are: ComputeDriver, Controller, DB, LoadBalancer . For example, the following command recovers the nodes in the DB node group from the specified NFS backup file. <pre>viocli recover -r DB -dn vio_os_db_20150830215406 -nfs 10.146.29.123:/backups</pre>
<code>-n, --node NODE</code>	Mandatory unless <code>ROLE</code> is specified	Recovers a specified node. You can specify multiple nodes in one command. Use the VM name as it appears in the VMware Integrated OpenStack manager. To view the name, select VMware Integrated OpenStack > OpenStack Deployments > [Deployment Name] . For example, the following command recovers the specified database nodes (VIO-DB-0, VIO-DB-1, and VIO-DB-2) from the specified NFS backup file. <pre>viocli recover -n VIO-DB-0 VIO-DB-1 VIO-DB-2 -dn vio_os_db_20150830215406 -nfs 10.146.29.123:/backups</pre>
<code>-dn, --dir-name BACKUP_NAME</code>	Mandatory for OpenStack database recovery	Indicates the time stamp label of the backup file to use to restore the database.
For database recovery, use one of the following positional arguments	Mandatory for OpenStack database recovery	<p>DIR_NAME Name of the NFS directory containing the database backup file.</p> <p>NFS_VOLUME Name or IP address of the target NFS volume and directory containing the database backup.</p> <p>Use the following format: <code>remote_host:remote_dir</code>. For example: <code>192.168.1.77:/backups</code>.</p>
<ul style="list-style-type: none"> ■ <code>DIR_NAME</code> ■ <code>NFS_VOLUME</code> 		

Parameter	Mandatory or Optional	Description
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

viocli restore Command

Use the `viocli restore` command to restore a deployment from a backup file previously created by using the `viocli backup` command. You can restore a backup of either management server data or of the OpenStack database.

The `viocli restore` command uses the following syntax.

```
viocli restore mgmt_server [-d [NAME]] <DIR_NAME | NFS_VOLUME> [-h] [-v]
viocli restore openstack_db [-d [NAME]] <DIR_NAME | NFS_VOLUME> [-h] [-v]
```

Parameter	Mandatory or Optional	Description
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
Use one of the following positional arguments	Mandatory	<i>DIR_NAME</i> Name of the NFS directory containing the backup file. <i>NFS_VOLUME</i> Name or IP address of the target NFS volume and directory in the format <i>remote_host:remote_dir</i> . For example: 192.168.1.77:/backups.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

The backup file of the VMware Integrated OpenStack management server is labeled with the time stamp `vio_ms_yyyymmddhhmmss`. The backup file of the VMware Integrated OpenStack database is labeled with the time stamp `vio_os_db_yyyymmddhhmmss`.

viocli rollback Command

Use the `viocli rollback` command to roll back a recent upgrade of VMware Integrated OpenStack. Use the VMware Integrated OpenStack manager in the vSphere Web Client to perform the rollback.

The `viocli rollback` command uses the following syntax.

```
viocli rollback [-d NAME] [-p] [-h] [-v] [-f]
```

Parameter	Mandatory or Optional	Description
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
-p, --progress	Optional	Show progress of the current upgrade operation.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.
-f, --force	Optional	Perform operation without requiring confirmation.

viocli services Command

Use the `viocli services` command to start or stop OpenStack services. The difference between `viocli deployment stop` and `viocli services stop` is that the former stops the entire cluster including virtual machines. The latter stops only the services running on the virtual machines in the cluster.

The `viocli services` command uses the following syntax.

```
viocli services ACTION [-d NAME] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
ACTION	Mandatory	start Start the deployment. stop Stop the deployment.
Use one of the following positional arguments:		
<ul style="list-style-type: none"> ■ start ■ stop 		
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

viocli show Command

Use the `viocli show` command to display a list of the nodes in a VMware Integrated OpenStack deployment, or to get detailed information about the deployment inventory.

The `viocli show` command uses the following syntax.

```
viocli show [-p] [-i] [-d NAME] [-h] [-v]
```

Parameter	Mandatory or Optional	Description
-p, --inventory-path	Optional	Displays the inventory path used for the current deployment .
-i, --inventory	Optional	Displays inventory file content used for the current deployment
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.

viocli upgrade Command

Use the `viocli upgrade` command to upgrade between major versions of VMware Integrated OpenStack. Use the VMware Integrated OpenStack manager in the vSphere Web Client to perform the upgrade.

The `viocli upgrade` command uses the following syntax.

```
viocli upgrade [-d NAME] [-n NEW_DEPLOYMENT_NAME] \
               [--public-vip PUBLIC_VIP] [--internal-vip INTERNAL_VIP] \
               [-p] [-h] [-v] [-f]
```

Parameter	Mandatory or Optional	Description
-d, --deployment <i>NAME</i>	Automatic	Name of deployment to use. Applied automatically. The default value is the name of the current deployment.
-n, --new-deployment <i>NEW_DEPLOYMENT_NAME</i>	Mandatory	Name of upgrade deployment.
--public-vip <i>PUBLIC_VIP</i>	Mandatory	Temporary public VIP address assigned to the new deployment.
--internal-vip <i>INTERNAL_VIP</i>	Mandatory	Temporary private VIP address assigned to the new deployment.
-p, --progress	Optional	Show progress of the current upgrade operation.
-h, --help	Optional	Show the use and arguments for this command.
-v, --verbose	Optional	Enter verbose mode.
-f, --force	Optional	Perform operation without requiring confirmation.

viocli volume-migrate Command

Use the `viocli volume-migrate` command to migrate one or more non-attached volumes from one datastore to another.

- Volume migration fails, if the volume is attached.

In this case, migrate the corresponding instance. Attached volumes migrate with the instance.

Note Corresponding volume shadow VMs do not migrate. To migrate such volume shadow VMs, run the [viocli ds-migrate-prep Command](#), then migrate the shadow VMs using the vSphere Web Client.

- Volume migration fails, if the volume has a storage policy that cannot be satisfied by the destination datastore.

You can force the migration by including the `--ignore-storage-policy` parameter. The command outputs a warning if the storage policy is ignored for the migration to a non-compliant datastore.

The `viocli volume-migrate` command uses the following syntax.

```
viocli volume-migrate [-d [NAME]] \
  [--source-dc [SRC_DC_NAME]] [--source-ds [SRC_DS_NAME]] \
  [--volume-ids [VOLUME_UUIDS]] [--ignore-storage-policy] \
  DEST_DC_NAME DEST_DS_NAME [-h] [-v]
```

Parameter	Mandatory or Optional	Description
<code>-d, --deployment NAME</code>	Automatic	Name of the deployment in which the volumes to be migrated. Applied automatically. The default value is the name of the current deployment.
<code>--source-dc SRC_DC_NAME</code>	Mandatory unless <code>VOLUME_UUIDS</code> is specified.	Identifies the source data center. Used with the <code>--source-ds</code> parameter uniquely to identify the datastore.
<code>--source-ds SRC_DS_NAME</code>	Mandatory unless <code>VOLUME_UUIDS</code> is specified.	Used with the <code>--source-dc</code> parameter uniquely to identify the datastore. For example, the following command migrates all the volumes from datastore DS-01 in data center DC-01 to datastore DS-02 in data center DC-02. <code>viocli volume-migrate --source-dc DC-01 --source-ds DS-01 DC-02 DS-02</code>
<code>--volume-ids VOLUME_UUIDS</code>	Mandatory unless <code>SRC_DC_NAME</code> and <code>SRC_DS_NAME</code> are specified.	Migrates one or more individual volumes specified by UUID value. To specify more than one volume, separate the UUIDs by commas. For example, the following command migrates two volumes specified by their UUID values to datastore DS-01 in data center DC-01. <code>viocli volume-migrate --volume-ids 25e121d9-1153-4d15-92f8-c92c10b4987f, 4f1120e1-9ed4-421a-b65b-908ab1c6bc50 DC-01 DS-01</code>

Parameter	Mandatory or Optional	Description
<code>--ignore-storage-policy</code>	Optional	Ignores storage policy compliance check. Include this parameter to prevent migration failure if the migrated volume includes a storage policy with which the destination datastore does not comply.
<code>DEST_DC_NAME</code>	Mandatory	Specifies the destination data center.
<code>DEST_DS_NAME</code>	Mandatory	Specifies the destination datastore.
<code>-h, --help</code>	Optional	Show the use and arguments for this command.
<code>-v, --verbose</code>	Optional	Enter verbose mode.

vRealize Automation Integration

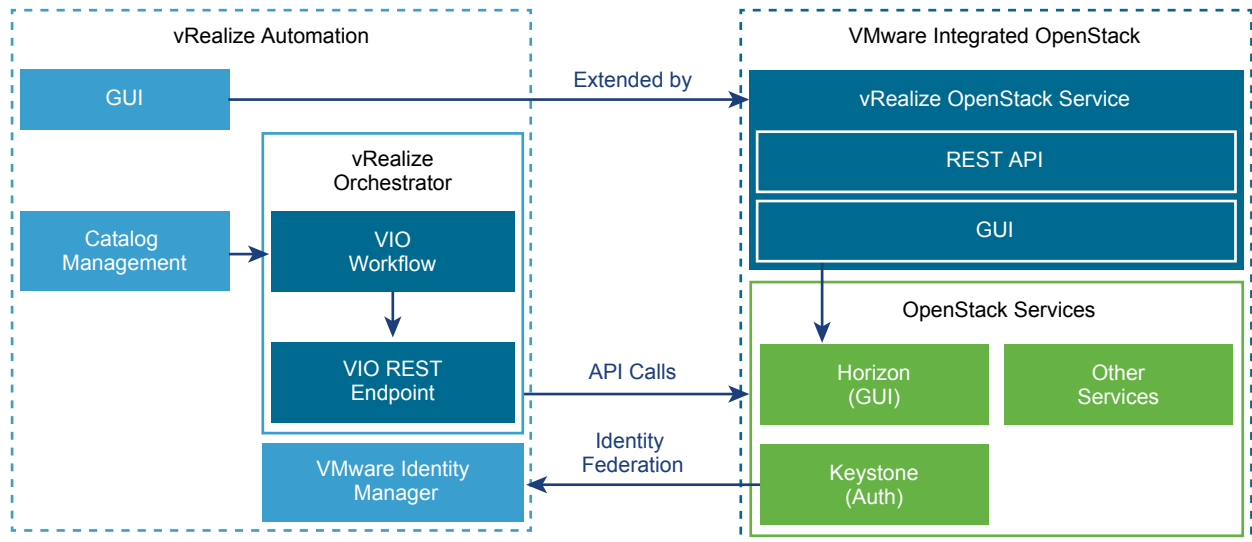
10

You can integrate VMware Integrated OpenStack with vRealize Automation and vRealize Orchestrator to enforce control and governance, manage OpenStack deployments as resource pools, and administrate VMware Integrated OpenStack from the vRealize Automation portal.

When you integrate VMware Integrated OpenStack with vRealize Automation, you can benefit from the following features:

- Securely use existing credentials to access cloud resources through integration with VMware Identity Manager.
- Manage all your OpenStack deployments from a single GUI through the new VMware Integrated OpenStack tab that appears in the vRealize Automation portal.
- Consume VMware Integrated OpenStack based infrastructure through vRealize Automation XaaS blueprints.
- Run OpenStack Heat workflows that provide on-demand network capabilities on OpenStack based resource pools
- Run workflows to manage VMs, projects, and networks.
- Create custom OpenStack workflows through the OpenStack API.

Figure 10-1. Integration Architecture



VMware Identity Manager Integration

By integrating VMware Integrated OpenStack with VMware Identity Manager you achieve a way to securely use existing credentials to access cloud resources such as servers, volumes, and databases, across multiple endpoints provided in multiple authorized clouds. You have a single set of credentials, without having to provision additional identities or log in multiple times. The credential is maintained by the user's Identity Provider.

Managing OpenStack Deployments Through the vRealize Automation Portal

If you have enabled the VMware Identity Manager integration, you can use the new VMware Integrated OpenStack tab that appears in the vRealize Automation portal. This tab embeds the VMware Integrated OpenStack dashboard in the vRealize Automation portal to allow for cloud administrators to manage OpenStack deployments from a single GUI. vRealize Automation administrator must enable the new tab and configure mappings to associate users to their respective projects. When a user who is associated to a project logs in to the vRealize Automation portal, they see the VIO tab.

vRealize Automation XaaS Blueprints Design

To consume vRealize Automation blueprints, you must install the vRealize Orchestrator Plug-in for OpenStack. vRealize Automation administrators can design and publish OpenStack blueprints. An approval chain and entitlement can also be configured. vRealize Automation users can request OpenStack catalog items that can be either approved or denied by users with assigned approval role.

vRealize Orchestrator Workflows

After you design vRealize Automation XaaS Blueprints, you consume them through the vRealize Orchestrator workflows that allow cloud administrators to automate user on-boarding and application deployment to OpenStack.

This section includes the following topics:

- [Integrate VMware Integrated OpenStack with vRealize Automation](#)
- [Designing and Publishing Blueprints](#)

Integrate VMware Integrated OpenStack with vRealize Automation

You integrate the two solutions by configuring the vRealize Automation tenant FQDN through the OpenStack management server and installing the vRealize Orchestrator OpenStack plug-in.

Procedure

- 1 Enable the service that integrates VMware Integrated OpenStack with vRealize Automation.
 - a Log in to the OpenStack management server.
 - b Run the command to configure the VMware Integrated OpenStack tab for your vRealize Automation tenant.

The `vra_tenant` value must be uppercase only.

```
viocli vros enable -d your_VIO_deployment_name --vra_tenant YOUR_VRA_TENANT --vra_host IP/FQDN
--vra_admin vRA_administrator_username --vros_host OMS_server_IP --verbose
```

- 2 Prior to being able to login a VMware Identity Manager user to VMware Integrated OpenStack, assign a role or project to the group that user belongs to.
 - a Create JSON file and configure mappings between the vIDM domain, the Openstack project, and the OpenStack group.

```
{
  "domain": {
    "name": "Default", //domain where VMware Identity Manager users will consume
    resource from.
  }
  "project": {
    "name": "fed-project1", // new OpenStack project for this group of users.
    "group": {
      "name": "fed-group1", // new OpenStack group for these vRealize Automation users
      "users": [ // list of vIDM users will be added to the group.
        {
          "name": "user1"
        },
        {
          "name": "user2"
        }
      ]
    }
  }
}
```

- b Apply the mappings configured in the JSON file.

```
vros_cli.sh add-mapping mappings_file.json
```

3 Deploy the vRealize Orchestrator OpenStack Plug-In.

- a Download the `o11n-openstack-plugin.2.0.0-XXX.vmoapp` file on your local file system from http://Your_OMS_server_IP/o11n-openstack-plugin/.

You can log in to either the vRealize Orchestrator embedded in your vRealize Automation server or the external vRealize Orchestrator connected to the vRealize Automation server.

- b Log in to the vRealize Orchestrator Control Center as `root`.
- c On the Control Center page, click **Manage Plug-Ins**.
- d In the Install Plug-In section, click **Browse**.

- e Select the `o11n-openstack-plugin.2.0.0-XXX.vmoapp` file from the local file system.

The vRealize Orchestrator server installs the plug-in. After the plug-in is installed, the Manage Plug-Ins page refreshes and indicates that the installation was successful. The page also displays a message prompting you to restart the vRealize Orchestrator server.

- f Click **Install**.
- g Click **Startup Options** to access controls and restart the vRealize Orchestrator server.

You can now administrate VMware Integrated OpenStack through the vRealize Automation portal, design, and consume blueprints.

Designing and Publishing Blueprints

You can design blueprints to run any of the available workflows. For example, you can design blueprints to create or delete OpenStack projects, and to deploy or delete Heat stacks.

Service Accounts for vRealize Orchestrator Workflows

You can use service accounts for authentication when you design blueprints. These accounts are used by default for authentication of the OpenStack REST API and users do not have to enter credentials every time they request services. When you use service accounts, set the credentials related parameters as not visible to the users. You create these accounts as standard OpenStack users.

Design an XaaS Blueprint for Creating or Deleting OpenStack Projects

You can design blueprints that create OpenStack projects when consumed.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > Deploy an OpenStack project** or **Delete an OpenStack project**, and click **Next**.

The blueprint form appears.

- 5 Complete the text boxes in the blueprint form.

Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the Project name parameter and the mandatory ones.

Parameter	Label	Value	Visibility to user
keystone_url	Keystone URL	Provide the URL to the Keystone Public Service.	False
admin_username	Admin Username	Provide a valid Active Directory or LDAP user name with administrative permissions.	False
admin_password	Admin Password	Provide a valid password for the Active Directory or LDAP user name.	False
admin_project_name	Admin Project Name	Enter the name for the administrative OpenStack project name where the OpenStack instances will be provisioned. The default is admin .	False
keystone_domain_name	Keystone Domain Name	Specify the Keystone domain for OpenStack user authentication. Keystone is the OpenStack Identity Service component. The default domain name is default.	False
new_project_name	New Project Name	Specify the name of the new OpenStack project where the OpenStack instances will be provisioned. You can leave this value empty, which enables the catalog user to specify the name.	True
user_name	Username	The catalog user's OpenStack account name with the Active Directory or LDAP domain suffix omitted.	True
quotas	quotas settings	Configure the predefined quota keys and values for the new project. <ul style="list-style-type: none"> ■ nova_instances = 120 ■ neutron_subnet = 130 ■ cinder_snapshots = 140 The quota key names use the <i>servicename_quotaname</i> format, where <i>servicename</i> is the name of the standard OpenStack service and <i>quotaname</i> is the name of the standard OpenStack service quota .	False

- 6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- 7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Deploying or Deleting Heat Stacks

You can design blueprints that deploy Heat stacks when consumed as a service.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > Deploy a Heat stack** or **Delete a Heat stack**, and click **Next**.
- 5 Select the template to be published and click **Next**.

The blueprint form appears.

- 6 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Label	Value	Visibility to user
stack_name	New Stack Name	Specify the name for the new Heat stack. You can leave this value empty to allow the catalog user to specify the new stack name.	True
template_content	Template Content	Provide the Heat template file content to be published. Configure as text area. A Heat template is a static architectural design of the orchestrated application, and are written in the HOT (Heat Orchestration Template) format.	True
environment_content	Environment Content	Provide the Heat environment file content. Configure as text area. The Heat environment file contains values for specific parameters in the Heat template. For authoring information, see the OpenStack documentation.	False
timeout	timeout in minutes	Specify the timeout period for this service in minutes.	False

7 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

8 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting VM Instances

You can design blueprints that create or delete VM instances.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create an instance or delete an instance**, and click **Next**.

The blueprint form appears.

- 5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description	Visibility to user
Instance name	Enter name for the VM.	True
Image name	Name of the VM image to clone from. Instance source is the template used to create an instance. You can use a snapshot of an existing instance, an image, or a volume, if enabled. You can also select to use persistent storage by creating a new volume.	True
Flavor name	Name of the flavor. Flavors manage the sizing for the compute, memory and storage capacity of the instance.	True
Network name	Enter name of an existing network. Networks provide the communication channels for instances in the cloud.	True

Parameter	Description	Visibility to user
Keypair name	Select an existing key pair, import a key pair, or generate a new key pair. A key pair allows you to connect over SSH to your newly created instance.	True
Availability zone	Optional. Defines the placement for allocation of virtual machines.	True
Timeout	Timeout for the workflow to wait for creation complete.	True

6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting an SSH Pair

You can design blueprints that create or delete SSH pairs.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- Log in to the vRealize Automation tenant as a service architect.
- Select **Design > XaaS > XaaS Blueprints**.
- Click **New**.
- Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create a keypair or delete a keypair**, and click **Next**.

The blueprint form appears.

5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints.

Parameter	Description of the Value	Visibility to Users
Keypair name	Name for the keypair	True

6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting a Glance Image

You can design blueprints that create OpenStack projects when consumed.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- Log in to the vRealize Automation tenant as a service architect.
- Select **Design > XaaS > XaaS Blueprints**.
- Click **New**.
- Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create an image or delete an image**, and click **Next**.

The blueprint form appears.

- Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Image name	Name of the new image.	True
Disk format	Format of the image file, for example: vmdk , qcow2 , etc.	True
Image location	URL to the image file, for example: https://ip:port/image1.vmdk	True
Is public	If you select true , image is shared with other projects, false restricts it to the current project.	True
Disk type	Enter the disk type, for example: streamOptimized or sparse .	True
Adapter type	Enter adapter type, for example lsiLogic .	True
Container format	Enter the container format, for example bare .	True
Timeout	Timeout in seconds for the workflow to wait for the image creation to finish.	True

- Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- 7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting a Network

You can design blueprints that create or delete provider networks.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create a network or delete a network**, and click **Next**.

The blueprint form appears.

- 5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
network name	Enter name for the network.	True
network type	Enter the type of the physical network. For example: flat , vlan , portgroup , or vxlan .	True
physical network	Enter the corresponding physical network for this new virtual network.	True
project name for this network	Enter the keystone project where this new virtual network is created.	True
is shared	If you select <code>true</code> , network is shared with other projects, <code>false</code> restricts it to the current project.	True
is admin state up	Select the state to start the network in.	True
segmentation_id	Enter the segmentation id of the physical network.	True
is external	Select whether this virtual network is an external network.	True

- 6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting a Subnet

You can design blueprints that create or delete subnets.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- Log in to the vRealize Automation tenant as a service architect.
- Select **Design > XaaS > XaaS Blueprints**.
- Click **New**.
- Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create a subnet or delete a subnet**, and click **Next**.

The blueprint form appears.

- Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Subnet name	Name for the subnet	True
Network name	Name of the network that the subnet is created under.	True
CIDR	IP address range in CIDR format.	True
DHCP	True or false according to your needs.	True
IP start	IP address starting allocation.	True
IP end	IP address ending allocation.	True
Gateway IP	IP address of the gateway.	True

- Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting a Floating IP

You can design blueprints that create or delete a floating IP address.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create a floating ip** or **delete a floating ip**, and click **Next**.

The blueprint form appears.

- 5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Floating IP	Enter the IP address of this floating IP.	True
External network name	Enter the name of the external network, where the floating IP is allocated from.	True

- 6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- 7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting a Logical Router

You can design blueprints that create or delete logical routers.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create a router or delete a router**, and click **Next**.

The blueprint form appears.

- 5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Router name	Enter name of the router.	True
External network	Enter name of the network where the router is uplinked.	True
Is admin state up	Select the state to start the router in.	True
Is distributed	Select whether the router is distributed or not.	True
Router type	Select between exclusive or shared.	True
Router size	Select between compact, large, or extra large.	True

- 6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- 7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Creating or Deleting a Security Group

You can design blueprints that create or delete a security group.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.

- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > create a security group** or **delete a security group**, and click **Next**.

The blueprint form appears.

- 5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Name	Name for the security group.	True
Description	(Optional) Enter a description for that group.	True

- 6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

- 7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Attaching or Detaching a Cinder Volume

You can design blueprints that attach or detach volumes to the server.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > attach volume to server** or **detach volume to server**, and click **Next**.

The blueprint form appears.

5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Instance name	Enter the name of the VM to attach to.	True
Volume name	Enter the name of the volume to be attached.	True

6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Design an XaaS Blueprint for Attaching or Detaching a Cinder Volume

You can design blueprints that attach or detach volumes to the server.

Prerequisites

- Verify that you have **service architect** user privileges to access the plug-in.
- Verify that the OpenStack Keystone service is connected to an Active Directory server for automatic synchronization. This configuration automatically synchronizes user accounts in Keystone

Procedure

- 1 Log in to the vRealize Automation tenant as a service architect.
- 2 Select **Design > XaaS > XaaS Blueprints**.
- 3 Click **New**.
- 4 Select the workflow **Orchestrator > Library > VMware Integrated OpenStack > attach volume to server** or **detach volume to server**, and click **Next**.

The blueprint form appears.

5 Complete the text boxes in the blueprint form.

You must enter custom labels for all parameters. Parameters listed below are specific for this blueprint. Keystone URL, Project name, Keystone domain name, User name, and Password parameters are mandatory for all blueprints. For the delete blueprint you only need the first parameter from the table and the mandatory ones.

Parameter	Description of Value	Visibility to Users
Instance name	Enter the name of the VM to attach to.	True
Volume name	Enter the name of the volume to be attached.	True

6 Click **Finish**.

The newly created blueprint appears in the list of XaaS blueprints.

7 To publish the blueprint, select it from the list and click **Publish**.

A blueprint becomes available for consumption only after you publish it.

Publish XaaS Blueprints as Catalog Items

After an XaaS blueprint is designed, you must publish it to the catalog to make it available to service catalog users. All catalog items must be associated with a service so that you can entitle users and groups to access the services. .

Prerequisites

Verify that you have **tenant administrator** user privileges to access the plug-in.

Procedure

1 Log in to the vRealize Automation tenant as a tenant administrator.

2 Select **Administration > Catalog Management > Services**.

3 Click the **New** icon.

4 Enter a name and description.

These values appear in the service catalog for the catalog users. For example, if it does not exist already, you can create the **OpenStack Services** category.

5 (Optional) To associate the OpenStack icon with the new service, click **Browse** and select the OpenStack icon.

6 Click **Finish**.

7 Select OpenStack Services and click **Manage Catalog Items** from the drop-down menu.

8 Click the **Add** icon and select the desired blueprints.

9 Add the selected blueprints to the current service and click **OK**.

10 Select OpenStack Services and click **Activate** from the drop-down menu.

11 (Optional) Select **Administration > Catalog Management > Catalog Items** and to update the added services with the OpenStack icon.

What to do next

You can now configure entitlements for users and groups to use the new OpenStack services. See the [vRealize Automation product documentation](#).

Request an OpenStack Service

Catalog users with the required entitlement can request the running of a configured service through the services catalog in vRealize Automation.

For more information about using entitlements and how they determine which users and groups can request specific catalog items or perform specific actions, see the [vRealize Automation product documentation](#).

Prerequisites

- Verify that you have the user privileges to access the plug-in as a **catalog user**.
- Verify that your user account possesses the required entitlement to access and request OpenStack Services.

Procedure

- 1 Log in to the vRealize Automation tenant as a catalog user.
- 2 Select **Catalog > OpenStack Services**.
- 3 Click **Request** for the service that you need.

The actual name of this service is determined by the person who creates the source blueprint.

The New Request page displays the parameters you need to supply.

- 4 On the **Request Information tab**, enter a brief description.
- 5 Click the **Step** tab and provide the required configurations.
- 6 Click **Submit**.

The request is submitted. After the request is approved, depending on the approval policy associated with the catalog service, the OpenStack service request runs.

- 7 To confirm if the request was successful, you can log in to OpenStack to verify that the OpenStack service request completed successfully.

You must log in to OpenStack as the same user who requested the new service.