VMware Integrated OpenStack Administration Guide

Update 2
Modified on 13 NOV 2018
VMware Integrated OpenStack 4.1

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

Contents

VMware Integrated OpenStack Administration Guide

1

The *VMware Integrated OpenStack Administration Guide* shows you how to perform administrative tasks in 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.

Terminology

For definitions of terms as they are used in this document, see the VMware Glossary at https://www.vmware.com/topics/glossary and the OpenStack Glossary at https://docs.openstack.org/doc-contrib-guide/common/glossary.html.

Updated Information

The *VMware Integrated OpenStack Administration Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *VMware Integrated OpenStack Administration Guide*.

Revision	Description
Update 2 (13 NOV 2018)	■ Added documents about creating availability zones. ■ Added document about creating provider security groups. ■ Various corrections and improvements
Update 1 (08 OCT 2018)	■ Added troubleshooting section ■ Updated command reference ■ Various corrections and improvements
13 JUL 2018	Updated Import Virtual Machines into VMware Integrated OpenStack to add a link to DCLI Cannot Connect to Server.
22 FEB 2018	Added document Configure MAC Learning.
18 JAN 2018	Initial release.

Deployment Configuration 2

You can modify the configuration of your VMware Integrated OpenStack deployment to add capacity, enable profiling, update credentials, and change or customize various settings.

This chapter includes the following topics:

Add Capacity to an OpenStack Deployment

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

Add Compute Clusters to Your Deployment

You can add compute clusters to your VMware Integrated OpenStack deployment to increase processing capacity.

Prerequisites

In vSphere, create the cluster that you want to add to your deployment.

If you want to add compute clusters from a separate compute vCenter Server instance, the following restrictions apply:

- You must deploy VMware Integrated OpenStack in HA mode with NSX-T Data Center networking. Other deployment and networking modes do not support adding compute clusters from separate vCenter Server instances.
- You cannot add compute clusters from separate compute vCenter Server instances in the same availability zone.

Procedure

1. In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.
2. Open the **Manage** tab.
3. If you want to add compute clusters from a separate vCenter Server instance, first add the instance to your deployment.
   a. Select the **Compute vCenter Server** tab.
   b. Click the **Add** (plus sign) icon at the top left of the pane.
   c. Enter the FQDN of the vCenter Server instance and administrator credentials and click **OK**.
4. Select the **Nova Compute** tab.
5. Click the **Add** (plus sign) icon at the top left of the pane.
6. Select the vCenter Server instance and availability zone for the compute cluster that you want to add and click **Next**.
7. Select the new compute cluster and click **Next**.

   The cluster you select must contain at least one host.
8. Select one or more datastores for the compute cluster to consume and click **Next**.
9. Select the management virtual machine and desired datastore and click **Next**.
10. Review the proposed configuration and click **Finish**.

The processing capacity of your deployment increases accordingly with the size of the additional compute cluster.

Add Storage to a Compute Node

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

Adding a datastore causes the compute service to restart and might temporarily interrupt OpenStack services.

Procedure

1. In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.
2. Open the **Manage** tab.
3. Open the **Nova Storage** tab and click the **Add** (plus sign) icon at the top left of the pane.
4. Select the cluster to which you want to add a datastore and click **Next**.
5. 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 in your VMware Integrated OpenStack deployment.

Adding a datastore causes the image service to restart and might temporarily interrupt OpenStack services.

Procedure

1. In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.
2. Open the **Manage** tab.
3. Open the **Glance Storage** tab and click the **Add** (plus sign) icon at the top left of the pane.
4. Select one or more datastores to add and click **Next**.
5. Review the proposed configuration and click **Finish**.

The storage capacity for the image service increases accordingly with the size of the additional datastore.

Modify Management and API Access Networks

You can add IP address ranges and DNS servers to the management network and API access network in your deployment.

Procedure

1. In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.
2. Open the **Manage** tab.
3. On the **Networks** tab, right-click the network that you want to modify.
   - To add an IP address range, select **Add IP Range** and enter the IP address range that you want to add to the network.

     You can click **Add IP Range** to add multiple IP address ranges at once.

     Important The management and API access networks cannot include more than 100 IP addresses each.
   - To change the DNS server, select **Change DNS** and enter the DNS servers that you want to add to the network.

     Modifying the DNS settings will briefly interrupt the network connection.

Update Component Credentials

Your VMware Integrated OpenStack deployment includes credentials that allow OpenStack to access and connect with your LDAP server, NSX Manager, and vCenter Server instance. You can modify these credentials in the VMware Integrated OpenStack vApp.

Important If you want to change the NSX-T Data Center password, perform the following steps:

1 Log in to a controller node and run the `systemctl stop neutron-server` command to stop the Neutron server service.

2 Change the password in NSX-T Data Center.

3 Change the password in VMware Integrated OpenStack as described in the following section.

The Neutron server service will restart after you change the password in VMware Integrated OpenStack.

Procedure

1 In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.

2 Open the **Manage** tab.

3 On the **Settings** tab, click **Change Password**.

The **Change Passwords** panel contains text boxes for updating the current LDAP server, NSX Manager, and vCenter Server credentials.

4 Enter the updated credentials and click **Submit**.

To retain the original settings for a component, leave the text boxes blank.

Add Certificates to Your Deployment

You can add digital certificates to your deployment in the VMware Integrated OpenStack vApp.

The certificates that you add must be signed by a certificate authority (CA) and created from a certificate signing request (CSR) generated by VMware Integrated OpenStack. Using wildcard certificates is not supported.

Procedure

1 In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.

2 Open the **Manage** tab.

3 On the **Settings** tab, select **OpenStack SSL Certificate**.

4 If you require a new CA-signed certificate, enter the information for the CSR and click **Generate**.

5 After you have obtained the certificate from the CA, click **Import** and select the certificate file.

The certificate is added to your deployment.

Configure Public API Rate Limiting

Limiting the rate of calls made to API services can make operations more reliable and reduce the incidence of orphaned objects during high load.

If a client exceeds the rate limit, it receives an HTTP `429 Too Many Requests` response. The `Retry-After` header in the response indicates how long the client must wait before making further calls.

You can enable rate limiting by service. For example, you might want to throttle Nova API service calls more tightly than Neutron API service calls.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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. Uncomment the `haproxy_throttle_period` parameter and set it to the number of seconds that clients must wait if a rate limit is exceeded.
5. If you want to configure rate limits for specific APIs, uncomment the `max_requests` and `request_period` parameters for those services and configure them as desired.

   The APIs that can be rate limited and the corresponding parameters are listed as follows.

Option	Description
haproxy_keystone_max_requests **haproxy_keystone_request_period**	Keystone API
haproxy_keystone_admin_max_requests **haproxy_keystone_admin_request_period**	Keystone administrator API
haproxy_glance_max_requests **haproxy_glance_request_period**	Glance API
haproxy_nova_max_requests **haproxy_nova_request_period**	Nova API
haproxy_nova_placement_max_requests **haproxy_nova_placement_request_period**	Nova placement API
haproxy_cinder_max_requests **haproxy_cinder_request_period**	Cinder API

Option	Description
haproxy_designate_max_requests **haproxy_designate_request_period**	Designate API
haproxy_neutron_max_requests **haproxy_neutron_request_period**	Neutron API
haproxy_heat_max_requests **haproxy_heat_request_period**	Heat API
haproxy_heat_cfn_max_requests **haproxy_heat_cfn_request_period**	Heat CloudFormation API
haproxy_heat_cloudwatch_max_requests **haproxy_heat_cloudwatch_request_period**	Heat CloudWatch API
haproxy_ceilometer_max_requests **haproxy_ceilometer_request_period**	Ceilometer API
haproxy_aodh_max_requests **haproxy_aodh_request_period**	Aodh API
haproxy_panko_max_requests **haproxy_panko_request_period**	Panko API

6 Deploy the updated configuration.

```
sudo viocli deployment configure --limit lb
```

Deploying the configuration briefly interrupts OpenStack services.

Example: Limiting Calls to the Neutron Public API

The following configuration limits calls to the Neutron public API. If a single source IP address sends more than 50 requests to the Neutron public API in a 10 second period, the load balancers will return HTTP 429 errors to all subsequent requests from that source address for a period of 60 seconds. After 60 seconds have passed, the source address can resume sending requests to the Neutron public API.

```
haproxy_throttle_period: 60
haproxy_neutron_max_requests: 50
haproxy_neutron_request_period: 10
```

Create a Custom Theme for the VMware Integrated OpenStack Dashboard

You can modify the styling, logos, and bookmark icon of the VMware theme on the VMware Integrated OpenStack dashboard.

You can upload a custom logo, style sheet, or bookmark icon to the OpenStack Management Server and configure it to display as your theme.

Prerequisites

- Custom logos should be 216 pixels long by 35 pixels wide. Graphics with different dimensions might not be displayed properly.
- Custom logo files must be in SVG or PNG format.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Create the `/opt/vmware/vio/custom/horizon` directory and transfer the desired theme files to this directory.
3. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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.
5. Configure the desired theme elements.
   - To configure a bookmark icon, uncomment the `horizon_favicon` parameter and set its value to the path of your icon file.
   - To configure a dashboard logo (displayed in the top-left corner of each page), uncomment the `horizon_logo` parameter and set its value to the path of your dashboard logo file.
   - To configure a login logo, uncomment the `horizon_logo_splash` parameter and set its value to the path of your login logo file.
   - To configure additional styles, uncomment the `horizon_custom_stylesheet` parameter and set its value to the path of your stylesheet.
   - To configure color code definitions, uncomment the `horizon_custom_variables` parameter and set its value to the path of your stylesheet.
6. Deploy the updated configuration.

   ```
   sudo viocli deployment configure
   ```

   Deploying the configuration briefly interrupts OpenStack services.

The VMware Integrated OpenStack dashboard loads your custom theme files. After the service becomes available, you can log in and choose the VMware theme to display your customizations.

Note If you edit or disable the custom theme at a later time, clear the browser cache so that the updated theme can be displayed.

Profile OpenStack Services

You can use OSProfiler to enable tracing for the core services in your OpenStack deployment. Tracing captures the response time of all API, RPC, driver, and database calls that are part of an OpenStack operation.

VMware Integrated OpenStack supports the profiling of Cinder, Glance, Heat, Neutron, and Nova commands. You can store profiler trace data with Ceilometer or vRealize Log Insight.

Prerequisites

- If you want to use Ceilometer to store trace data, enable Ceilometer. See "Enable the Ceilometer Component" in the *VMware Integrated OpenStack Installation and Configuration Guide*.
- If you want to use vRealize Log Insight to store trace data, deploy and configure vRealize Log Insight. See the *Getting Started* document for vRealize Log Insight.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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. Uncomment the `os_profiler_enabled` parameter and set its value to **`true`**.
5. Uncomment the `os_profiler_hmac_keys` parameter and enter a password for OSProfiler.
6. If you are using vRealize Log Insight, uncomment the `os_profiler_connection_string` parameter and set its value to the location of your vRealize Log Insight server.

   Enter the vRealize Log Insight server address in the following format: "`loginsight://`*`username`*`:`*`password`*`@`*`loginsight-ip`*"

   Specify the user name and password of a user with the `USER` role on your vRealize Log Insight deployment.
7. Deploy the updated configuration.

   ```
   sudo viocli deployment configure
   ```

   Deploying the configuration briefly interrupts OpenStack services.
8. If you are using vRealize Log Insight, log in to the controller node and set the `OSPROFILER_CONNECTION_STRING` environment variable to the vRealize Log Insight server address that you specified in the `custom.yml` file.

   ```
   export OSPROFILER_CONNECTION_STRING="loginsight://username:password@loginsight-ip"
   ```

You can now enable profiling on OpenStack commands. Run the desired command with the `--profile` parameter and specify your OSProfiler password. The command outputs a profiling trace UUID. Run OSProfiler with that UUID to generate a report. The following example profiles the `cinder list` command:

```
cinder --profile osprofiler-password list
osprofiler trace show --html profiling-uuid
```

Create a Tenant Virtual Data Center

You can create tenant virtual data centers to enable secure multi-tenancy and resource allocation. These data centers can be created on different compute nodes that offer specific service level agreements for each telecommunication workload.

Important This feature is offered in VMware Integrated OpenStack Carrier Edition only. For more information, see "VMware Integrated OpenStack Licensing" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

Project quotas limit OpenStack resources across multiple compute nodes or availability zones, but they do not guarantee resource availability. By creating a tenant virtual data center to allocate CPU and memory for an OpenStack project on a compute node, you provide a resource guarantee for tenants and avoid noisy neighbor scenarios in a multi-tenant environment.

The tenant virtual data center allocates resources at the compute node level. You can also allocate resources on the virtual network function (VNF) level using the same flavor. For instructions, see Configure QoS Resource Allocation for Instances.

Tenant virtual data centers are managed using the `viocli` utility. For more information, see viocli inventory-admin Command.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Create a tenant virtual data center.

```
viocli inventory-admin create-tenant-vdc --project-id project-uuid --compute compute-node --name
display-name [--cpu-limit max-cpu-mhz] [--cpu-reserve min-cpu-mhz] [--mem-limit max-memory-mb] [--
mem-reserve min-memory-mb]
```

3 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.

4 Select the **admin** project from the drop-down menu in the title bar.

5 Configure a flavor to use the tenant virtual data center.

   a Select **Admin > System > Flavors**.

   b Create a new flavor or choose an existing flavor to use for passthrough.

   c Select **Update Metadata** next to the flavor that you want to use.

d In the **Available Metadata** pane, expand **VMware Policies** and click the **Add** (plus sign) icon next to **Tenant Virtual Datacenter**.

e Set the value of `vmware:tenant_vdc` to the UUID of the tenant virtual data center and click **Save**.

You can run the `viocli inventory-admin list-tenant-vdcs` command on the OpenStack Management Server to find the UUID of all tenant virtual data centers.

The tenant virtual data center is created. You can now launch instances in the tenant virtual data center by configuring them with the flavor that you modified in this procedure.

What to do next

You can display the resource pools in a tenant virtual data center by running the `viocli inventory-admin show-tenant-vdc --id` *`tvdc-uuid`* command. Each resource pool is listed with its provider ID, project ID, status, minimum and maximum CPU, minimum and maximum memory, and compute node information. If a tenant virtual data center includes multiple resource pools, the first row displays aggregate information for all pools.

You can update your tenant virtual data centers by running the `viocli inventory-admin update-tenant-vdc` command. For specific parameters, see viocli inventory-admin Command.

You can delete an unneeded tenant virtual data center by running the `viocli inventory-admin delete-tenant-vdc --id` *`tvdc-uuid`* command.

Neutron Network Configuration 3

You can create provider and external networks for your VMware Integrated OpenStack deployment, configure availability zones, and perform other advanced networking tasks.

This chapter includes the following topics:

Create a Provider Network

Provider networks map to physical networks in your data center, and their networking functions are performed by physical devices.

A provider network can be dedicated to one project or shared among multiple projects. Tenants can create virtual machines in provider networks or connect their tenant networks to a provider network through a Neutron router.

The specific configuration for creating a provider network depends on the networking mode of your VMware Integrated OpenStack deployment.

Create a Provider Network with NSX-T Data Center

With NSX-T Data Center networking, you can create a VLAN-based provider network.

Prerequisites

- Define a VLAN for the provider network and record its ID.
- To use DHCP with VM form-factor NSX Edge nodes, enable forged transmit and promiscuous mode on the port group containing the edge nodes. For instructions, see "Configure the Security Policy for a Distributed Port Group or Distributed Port" in the *vSphere Networking* document.

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 > Network > Networks**.
4. Click **Create Network** and configure the provider network.

Option	Description
Name	Enter a name for the network.
Project	Select the desired project from the drop-down menu.
Provider Network Type	Select **VLAN** from the drop-down menu.
Physical Network	Enter the UUID of the VLAN transport zone.
Segmentation ID	Enter the VLAN ID defined for the provider network.

5. Select **Enable Admin State** and **Create Subnet**.
6. If you want multiple projects to use the provider network, select **Shared**.
7. Click **Next** and configure the subnet.

Option	Description
Subnet Name	Enter a name for the subnet.
Network Address	Enter the IP address range for the subnet in CIDR format (for example, 192.0.2.0/24).
IP Version	Select **IPv4** or **IPv6**.
Gateway IP	Enter the gateway IP address. If you do not enter a value, the first IP address in the subnet is used. If you do not want a gateway on the subnet, select **Disable Gateway**.

8 (Optional) Configure additional settings for the subnet.

   a Under **Allocation Pools**, enter IP address pools from which to allocate the IP addresses of virtual machines created on the network. Enter pools as two IP addresses separated by a comma (for example, `192.0.2.10,192.0.2.15`). If you do not specify any IP address pools, the entire subnet is available for allocation.

   b Under **DNS Name Servers**, enter the IP address of one or more DNS servers to use on the subnet.

   c Under **Host Routes**, enter additional routes to advertise to the hosts on the subnet. Enter routes as the destination IP address in CIDR format and the next hop separated by a comma (for example, `192.0.2.0/24,192.51.100.1`).

9 Click **Create**.

Create a Provider Network with NSX Data Center for vSphere

With NSX Data Center for vSphere networking, you can create a flat, VLAN-based, port group-based, or VXLAN-based provider network.

Prerequisites

- If you want to create a VLAN-based network, define a VLAN for the provider network and record its ID.
- To use DHCP on a VLAN-based network with VM form-factor NSX Edge nodes, you must enable forged transmit and promiscuous mode on the port group containing the edge nodes. For instructions, see "Configure the Security Policy for a Distributed Port Group or Distributed Port" in the *vSphere Networking* document.
- If you want to create a port group-based network, create a port group for the provider network and record its managed object identifier (MOID).

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 > Network > Networks**.

4 Click **Create Network** and configure the provider network.

Option	Description
Name	Enter a name for the network.
Project	Select the desired project from the drop-down menu.
Provider Network Type	Select **Flat**, **VLAN**, **Port Group**, or **VXLAN** from the drop-down menu.

Option	Description
Physical Network	■ If you selected **Flat** or **VLAN** for the network type, enter the MOID of the distributed switch for the provider network. ■ If you selected **Port Group** for the network type, enter the MOID of the port group for the provider network. ■ If you selected **VXLAN** for the network type, this value is determined automatically.
Segmentation ID	If you selected **VLAN** for the network type, enter the VLAN ID defined for the provider network.

5 Select **Enable Admin State** and **Create Subnet**.

6 If you want multiple projects to use the provider network, select **Shared**.

7 Click **Next** and configure the subnet.

Option	Description
Subnet Name	Enter a name for the subnet.
Network Address	Enter the IP address range for the subnet in CIDR format (for example, 192.0.2.0/24).
IP Version	Select **IPv4** or **IPv6**.
Gateway IP	Enter the gateway IP address. If you do not enter a value, the first IP address in the subnet is used. If you do not want a gateway on the subnet, select **Disable Gateway**.

8 (Optional) Configure additional settings for the subnet.

a Under **Allocation Pools**, enter IP address pools from which to allocate the IP addresses of virtual machines created on the network. Enter pools as two IP addresses separated by a comma (for example, `192.0.2.10,192.0.2.15`). If you do not specify any IP address pools, the entire subnet is available for allocation.

b Under **DNS Name Servers**, enter the IP address of one or more DNS servers to use on the subnet.

c Under **Host Routes**, enter additional routes to advertise to the hosts on the subnet. Enter routes as the destination IP address in CIDR format and the next hop separated by a comma (for example, `192.0.2.0/24,192.51.100.1`).

9 Click **Create**.

Create a Provider Network with VDS Networking

With VDS networking, you can create a VLAN-based provider network.

Prerequisites

Define a VLAN for the provider network and record its ID.

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 > Network > Networks**.
4. Click **Create Network** and configure the provider network.

Option	Description
Name	Enter a name for the network.
Project	Select the desired project from the drop-down menu.
Provider Network Type	Select **VLAN** from the drop-down menu.
Physical Network	Enter **dvs**.
Segmentation ID	Enter the VLAN ID defined for the provider network.

5. Select **Enable Admin State** and **Create Subnet**.
6. If you want multiple projects to use the provider network, select **Shared**.
7. Click **Next** and configure the subnet.

Option	Description
Subnet Name	Enter a name for the subnet.
Network Address	Enter the IP address range for the subnet in CIDR format (for example, 192.0.2.0/24).
IP Version	Select **IPv4** or **IPv6**.
Gateway IP	Enter the gateway IP address. If you do not enter a value, the first IP address in the subnet is used. If you do not want a gateway on the subnet, select **Disable Gateway**.

8. (Optional) Configure additional settings for the subnet.
   a. Under **Allocation Pools**, enter IP address pools from which to allocate the IP addresses of virtual machines created on the network. Enter pools as two IP addresses separated by a comma (for example, **192.0.2.10,192.0.2.15**). If you do not specify any IP address pools, the entire subnet is available for allocation.
   b. Under **DNS Name Servers**, enter the IP address of one or more DNS servers to use on the subnet.
   c. Under **Host Routes**, enter additional routes to advertise to the hosts on the subnet. Enter routes as the destination IP address in CIDR format and the next hop separated by a comma (for example, **192.0.2.0/24,192.51.100.1**).
9. Click **Create**.

Create an External Network

External networks act as floating IP address pools to provide external access for instances in your deployment.

An external network can be dedicated to one project or shared among multiple projects. Tenants cannot create virtual machines in external networks.

The specific configuration for creating an external network depends on the networking mode of your VMware Integrated OpenStack deployment.

Create an External Network with NSX-T Data Center

For NSX-T Data Center deployments, you create an external network to contain the floating IP addresses of future tenant logical (tier-1) routers.

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 > Network > Networks**.

4 Click **Create Network** and configure the provider network.

Option	Description
Name	Enter a name for the network.
Project	Select the desired project from the drop-down menu.
Provider Network Type	Select **Local** to connect tenant logical routers to the default tier-0 router or **External** to connect tenant logical routers to another tier-0 router.
Physical Network	If you selected **External** as the provider network type, enter the UUID of the tier-0 router to which you want to connect future tenant logical routers.

5 Select **Enable Admin State**, **External Network**, and **Create Subnet**.

6 If you want multiple projects to use the external network, select **Shared**.

7 Click **Next** and configure the subnet.

Option	Description
Subnet Name	Enter a name for the subnet.
Network Address	Enter the IP address range for the subnet in CIDR format (for example, 192.0.2.0/24).
IP Version	Select **IPv4** or **IPv6**.
Gateway IP	Enter the gateway IP address. If you do not enter a value, the first IP address in the subnet is used. If you do not want a gateway on the subnet, select **Disable Gateway**.

8 Click **Next** and deselect **Enable DHCP**.

9 (Optional) Configure additional settings for the subnet.

   a Under **Allocation Pools**, enter IP address pools from which to allocate the floating IP addresses of tenant logical routers. Enter pools as two IP addresses separated by a comma (for example, `192.0.2.10,192.0.2.15`). If you do not specify any IP address pools, the entire subnet is available for allocation.

   b Under **DNS Name Servers**, enter the IP address of one or more DNS servers to use on the subnet.

   c Under **Host Routes**, enter additional routes to advertise to the hosts on the subnet. Enter routes as the destination IP address in CIDR format and the next hop separated by a comma (for example, `192.0.2.0/24,192.51.100.1`).

10 Click **Create**.

Create an External Network with NSX Data Center for vSphere

With NSX Data Center for vSphere networking, you can create a flat, VLAN-based, port group-based, or VXLAN-based external network.

Prerequisites

- If you want to create a VLAN-based network, define a VLAN for the external network and record its ID.
- If you want to create a port group-based network, create a port group for the external network and record its managed object identifier (MOID).

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 > Network > Networks**.

4 Click **Create Network** and configure the provider network.

Option	Description
Name	Enter a name for the network.
Project	Select the desired project from the drop-down menu.
Provider Network Type	Select **Flat**, **VLAN**, **Port Group**, or **VXLAN** from the drop-down menu.
Physical Network	■ If you selected **Flat** or **VLAN** for the network type, enter the MOID of the distributed switch for the provider network. ■ If you selected **Port Group** for the network type, enter the MOID of the port group for the provider network. ■ If you selected **VXLAN** for the network type, this value is determined automatically.
Segmentation ID	If you selected **VLAN** for the network type, enter the VLAN ID defined for the provider network.

5 Select **Enable Admin State**, **External Network**, and **Create Subnet**.

6 If you want multiple projects to use the provider network, select **Shared**.

7 Click **Next** and configure the subnet.

Option	Description
Subnet Name	Enter a name for the subnet.
Network Address	Enter the IP address range for the subnet in CIDR format (for example, 192.0.2.0/24).
IP Version	Select **IPv4** or **IPv6**.
Gateway IP	Enter the gateway IP address. If you do not enter a value, the first IP address in the subnet is used. If you do not want a gateway on the subnet, select **Disable Gateway**.

8 Click **Next** and deselect **Enable DHCP**.

9 (Optional) Configure additional settings for the subnet.

   a Under **Allocation Pools**, enter IP address pools from which to allocate the floating IP addresses of tenant logical routers. Enter pools as two IP addresses separated by a comma (for example, **`192.0.2.10,192.0.2.15`**). If you do not specify any IP address pools, the entire subnet is available for allocation.

   b Under **DNS Name Servers**, enter the IP address of one or more DNS servers to use on the subnet.

   c Under **Host Routes**, enter additional routes to advertise to the hosts on the subnet. Enter routes as the destination IP address in CIDR format and the next hop separated by a comma (for example, **`192.0.2.0/24,192.51.100.1`**).

10 Click **Create**.

Create a Layer 2 Bridge

A Layer 2 bridge allows compute nodes on an overlay network to communicate with a physical VLAN.

Create a Layer 2 Bridge with NSX-T Data Center

You can create a Layer 2 bridge in NSX-T Data Center through a bridge cluster.

Prerequisites

In NSX-T Data Center, create a bridge cluster that includes two dedicated ESXi hosts. See "Create a Bridge Cluster" in the *NSX-T Administration Guide*.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Log in to the controller node as `viouser`.

3 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

4 Create a logical Layer 2 gateway, specifying the UUID of the NSX-T Data Center bridge cluster as the device name.

```
neutron l2-gateway-create gateway-name --device name=bridge-cluster-uuid,interface_names="temp"
```

The interface name value is ignored, and the name is automatically assigned.

5 Create the logical Layer 2 gateway connection using the gateway created in the previous step.

```
neutron l2-gateway-connection-create gateway-name network-name --default-segmentation-id=vlan-id
```

Compute nodes on the overlay network can now access the specified VLAN.

Create a Layer 2 Bridge with NSX Data Center for vSphere

You can create a Layer 2 bridge in NSX Data Center for vSphere through a port group.

Prerequisites

Create a port group and tag it with the ID of the VLAN to which you want to connect your compute nodes.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Log in to the controller node as `viouser`.

3 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

4 Create a logical Layer 2 gateway, specifying the managed object identifier (MOID) of the port group as the interface name.

```
neutron l2-gateway-create gateway-name --device name=temp,interface_names="portgroup-moid"
```

NSX Data Center for vSphere creates a dedicated distributed logical router (DLR) from the backup edge pool. The device name value is ignored, and the object is automatically assigned a name in the format "L2 bridging-*gateway-id*".

5 Create the logical Layer 2 gateway connection using the gateway created in the previous step.

```
neutron l2-gateway-connection-create gateway-name network-name --default-segmentation-id=vlan-id
```

VXLAN compute nodes can now access the specified VLAN.

Configure VLAN Transparency

VLAN-transparent networks allow tagged packets to pass through without tags being removed or changed.

Note For VDS deployments, only provider networks can be transparent. For NSX Data Center for vSphere and NSX-T Data Center deployments, only tenant networks can be transparent.

To enable VLAN transparency on a network, include the `--transparent-vlan` parameter and disable port security when you create the network. For example:

```
openstack network create network-name --project project-uuid --transparent-vlan --disable-port-security
```

Configure MAC Learning

MAC learning enables network connectivity for multiple MAC addresses behind a single vNIC. MAC learning is useful for distributing workloads in large OpenStack deployments.

MAC learning in VMware Integrated OpenStack is implemented differently for NSX-T Data Center and NSX Data Center for vSphere deployments.

- For NSX-T Data Center deployments, MAC learning in VMware Integrated OpenStack is provided by NSX-T Data Center MAC learning. For more information, see "Understanding MAC Management Switching Profile" in the *NSX-T Administration Guide*.
- For NSX Data Center for vSphere deployments, MAC learning in VMware Integrated OpenStack is implemented by enabling forged transmit and promiscuous mode. The guest must request promiscuous mode.

The following conditions apply to MAC learning:

- MAC learning is not compatible with port security or security groups.
- For NSX Data Center for vSphere deployments, performance will be affected because vNICs that request promiscuous mode receive a copy of every packet.
- For NSX Data Center for vSphere deployments, no RARP requests are generated for the multiple MAC addresses behind a single vNIC when a virtual machine is migrated with vMotion. This can result in a loss of connectivity.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Switch to the `root` user and load the cloud administrator credentials file.

   ```
   sudo su -
   source ~/cloudadmin.rc
   ```

3 Disable port security and security groups on the port where you want to configure MAC learning.

```
neutron port-update port-uuid --port-security-enabled false --no-security-groups
```

4 Enable MAC learning on the port.

```
neutron port-update port-uuid --mac-learning-enabled true
```

Create a Neutron Availability Zone with NSX Data Center for vSphere

Availability zones enable high availability for network resources. You can place nodes running the same service in different availability zones to ensure that service is not interrupted in the event of a failure in one zone.

Prerequisites

- Create an edge cluster for the new availability zone.
- Create a resource pool on the new edge cluster.
- Configure the new edge cluster to use the appropriate distributed switch. You can create a new distributed switch for the zone if desired.
- In NSX Data Center for vSphere, create a transport zone that includes the new edge cluster.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

```
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 Uncomment the `nsxv_availability_zones` parameter and set its value to the name of the availability zone that you want to create.

The value of this parameter can include multiple availability zones. Separate multiple names with commas (,).

5 Uncomment the `nsxv_availability_zones_detail` parameter and configure it for your new availability zone.

Option	Description
zone_name	Enter the name of the availability zone that you want to configure.
resource_pool_id	Enter the managed object identifier (MOID) of the resource pool that you created for the new availability zone.

Option	Description
datastore_id	Enter the MOID of the datastore that you want to use for the new availability zone.
edge_ha	Enter **True** to enable high availability for edge nodes or **False** to disable it.
ha_datastore_id	Enter the MOID of the datastore that you want to use for high availability for edge nodes. If you set `edge_ha` to `False`, do not specify a value for the `ha_datastore_id` parameter.
external_network	Enter the MOID of the external network port group on the distributed switch for the new availability zone.
vdn_scope_id	Enter the MOID of the transport zone that you created for the new availability zone.
mgt_net_id	Enter the MOID of the management network for your deployment.
mgt_net_proxy_ips	Enter the IP addresses of the metadata proxy server for your deployment.
dvs_id	Enter the MOID of the distributed switch for the new availability zone.

Ensure that there is one copy of the preceding parameters for each availability zone configured.

6 Deploy the updated configuration.

```
sudo viocli deployment configure
```

Deploying the configuration briefly interrupts OpenStack services.

What to do next

To specify an availability zone for a network, include the `--availability-zone-hint` *az-name* parameter when creating the network.

Manage Edge HA with NSX Data Center for vSphere

For NSX Data Center for vSphere deployments, you can enable HA for NSX Edge nodes and specify host groups in which the nodes will be placed.

Prerequisites

- Verify that your edge cluster has at least two hosts. If not, you might receive an anti-affinity error.
- If you want to specify edge host groups, create and configure the host groups in vSphere.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

```
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 Uncomment the `nsxv_edge_ha` parameter and set its value to **True**.

5 If you want to use edge host groups, uncomment the `nsxv_edge_host_groups` parameter and set its value to the two edge host groups that you created, separated by a comma (,).

6 Deploy the updated configuration.

```
sudo viocli deployment configure
```

Deploying the configuration briefly interrupts OpenStack services.

7 Log in to the controller node as `viouser`.

8 If you specified host groups, update your environment to include them.

```
sudo -u neutron nsxadmin -o nsx-update -r edges --property hostgroup=all
```

9 If your environment already includes NSX Edge nodes, enable HA on those nodes and migrate them to the specified host groups.

a Enable high availability on each existing NSX Edge node.

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

To find the ID of an NSX Edge node, you can run the `sudo -u neutron nsxadmin -r edges -o nsx-list` command.

b Migrate all existing edge nodes to the specified host groups.

```
sudo -u neutron nsxadmin -r edges -o nsx-update --property hostgroup=all
```

If you want to migrate only specific edge nodes, you can use the following command:

```
sudo -u neutron nsxadmin -o nsx-update -r edges -p edge-id=edge-node-id -p hostgroup=True
```

Edge HA is enabled for the desired nodes. If you specified edge host groups, current and future edge nodes are created in those groups.

What to do next

You can update the edge host groups in `custom.yml` after the original configuration. After deploying `custom.yml`, run the following commands to update the environment:

```
sudo -u neutron nsxadmin -o nsx-update -r edges --property hostgroup=clean
sudo -u neutron nsxadmin -o nsx-update -r edges --property hostgroup=all
```

Then perform Step 9 again to migrate edge nodes to the new host groups.

Specify Tenant Router Types for NSX Data Center for vSphere

For NSX Data Center for vSphere deployments, you can restrict the router types available to tenants and specify a default router type.

Note Administrators can create routers of any type.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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. Uncomment the `nsxv_tenant_router_types` parameter and specify the router types that you want to make available to tenants.

   You can enter **`exclusive`**, **`shared`**, **`distributed`**, or any combination separated by commas (,).

   The values of the `nsxv_tenant_router_types` parameter are used in order as the default router types.
5. Deploy the updated configuration.

   ```
   sudo viocli deployment configure
   ```

   Deploying the configuration briefly interrupts OpenStack services.

Tenants can create routers only of the types listed. If a tenant creates a router without specifying a type, the first available type is used by default.

Configure Dynamic Routing for Neutron Networks with NSX Data Center for vSphere

You can configure BGP dynamic routing for the provider and tenant networks in your environment.

After you enable BGP, the logical subnets created by your tenants are advertised outside of your environment without requiring source NAT or floating IP addresses. You must first create a VXLAN external network that you later use as internal interface for your gateway edges.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Log in to the controller node as `viouser`.

3 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

4 Create an IPv4 address scope for future tenant subnets and the subnet of your external VXLAN network.

```
neutron address-scope-create name 4
```

5 Create a subnet pool for the external network.

```
neutron subnetpool-create external-pool --pool-prefix network-address --default-prefixlen prefix-bits --address-scope scope-name --shared
```

Option	Description
external-pool	Enter a name for the subnet pool.
--pool-prefix	Enter the network address of the subnet pool in CIDR format (for example, 192.0.2.0/24). Subnets will be allocated from this network.
--default-prefixlen	Enter the network prefix length (in bits) to use for new subnets that are created without specifying a prefix length.
--address-scope	Enter the name of the IPv4 address scope that you created in Step 4.

6 Create a subnet pool for tenant networks.

```
neutron subnetpool-create tenant-pool --pool-prefix network-address --default-prefixlen prefix-bits --address-scope scope-name --shared
```

Note OpenStack will advertise this subnet pool to the physical fabric. Specify a prefix that is not currently in use.

Option	Description
tenant-pool	Enter a name for the subnet pool.
--pool-prefix	Enter the network address of the subnet pool in CIDR format (for example, 192.51.100.0/24). Subnets will be allocated from this network.
--default-prefixlen	Enter the network prefix length (in bits) to use for new subnets that are created without specifying a prefix length.
--address-scope	Enter the name of the IPv4 address scope that you created in Step 4.

7 Create a VXLAN-based external network.

```
neutron net-create network-name --provider:network_type vxlan --router:external
```

This command creates a new logical switch in NSX Data Center for vSphere.

8 Create a subnet on the external network.

The subnet must have DHCP disabled and no gateway.

```
neutron subnet-create external-network external-subnet-address --name external-subnet --
allocation-pool start=subnet-ip1,end=subnet-ip2 --subnetpool provider-subnet-pool --no-gateway --
disable-dhcp
```

Option	Description
external-network	Enter the name of the VXLAN-based external network that you created in Step 7.
external-subnet-address	Enter the network address for the subnet in CIDR format (for example, 192.51.100.0/28).
--name	Enter a name for the subnet.
--allocation-pool	Enter the first and last IP addresses of the range that you want to allocate from this subnet.
--subnetpool	Enter the subnet pool that you created in Step 5 for the external network.

9 Create BGP edge nodes.

```
sudo -u neutron nsxadmin -r bgp-gw-edge -o create --property name=edge-name --property local-
as=local-as-number --property external-iface=portgroup-moid:mgmt-network-ip --property internal-
iface=physical-net-id:external-network-ip
```

Option	Description
name	Enter a name for the BGP edge node.
local-as	Enter the local AS number for the edge node. The edges and physical routers cannot be in the same AS.
external-iface	Enter the managed object identifier (MOID) of the port group associated with the VLAN that connects the edge nodes to the physical routers. After the colon, enter the IP address of the edge node on the management network.
internal-iface	Enter the virtual wire identifier of the VXLAN-based external network. After the colon, enter the IP address of the edge node on the physical network. To find the virtual wire identifier, run the `openstack network show external-network-name` command and locate the value of the `provider:physical_network` parameter.

10 Enable BGP advertisement on the edge nodes.

```
sudo -u neutron nsxadmin -r routing-redistribution-rule -o create --property gw-edge-ids=edge1-
id,edge2-id --property learner-protocol=bgp --property learn-from=connected,bgp --property
action=permit
```

For the `gw-edge-ids` parameter, use the edge identifier (for example, `edge-4`) instead of the name. You can run the `sudo -u neutron nsxadmin -r bgp-gw-edge -o view` command to display the identifier of each BGP edge node.

11 Establish a BGP neighbor relationship between the edge nodes.

```
sudo -u neutron nsxadmin -r bgp-neighbour -o create --property gw-edge-ids=edge1-id,edge2-id --property ip-address=physical-router1-ip --property remote-as=remote-as-number --property password=bgp-password
```

Option	Description
gw-edge-ids	Enter the edge identifier of each node, separated by a comma.
ip-address	Enter the IP address on the physical router.
remote-as	Enter the AS number of the physical routers connected to the edge nodes.
password	Enter the BGP password.

12 Configure your physical routers.

a Ensure that the AS of the physical routers is the remote AS of the edge nodes.

b Configure the edge nodes as BGP neighbors.

c Set each router to advertise itself as a dynamic gateway.

13 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
```

14 (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.

Create a Provider Security Group

You can create a provider security group to block specific traffic for a project.

Standard security groups are created and managed by tenants, whereas provider security groups are created and managed by the cloud administrator. Provider security groups take precedence over standard security groups and are enforced on all virtual machines in a project.

For instructions about standard security groups, see "Working with Security Groups" in the *VMware Integrated OpenStack User's Guide*.

Procedure

1 Log in to the OpenStack Management Server.

2 Create a provider security group for a specific project.

```
neutron security-group-create group-name --provider=True --tenant-id=project-id
```

3 Create rules for the provider security group.

Note Provider security group rules block the specified traffic, whereas standard security rules allow the specified traffic.

```
neutron security-group-rule-create group-name --tenant-id=project-id [--description rule-description] [--direction {ingress | egress}] [--ethertype {IPv4 | IPv6}] [--protocol protocol] [--port-range-min range-start --port-range-max range-end] [--remote-ip-prefix ip/prefix | --remote-group-id remote-security-group]
```

Option	Description
group-name	Enter the provider security group created in Step 2.
--tenant-id	Enter the ID of the desired project.
--description	Enter a custom description of the rule.
--direction	Specify **`ingress`** to block incoming traffic or **`egress`** to block outgoing traffic. If you do not include this parameter, `ingress` is used by default.
--ethertype	Specify **`IPv4`** or **`IPv6`**. If you do not include this parameter, `IPv4` is used by default.
--protocol	Specify the protocol to block. Enter an integer representation ranging from 0 to 255 or one of the following values: ■ **`icmp`** ■ **`icmpv6`** ■ **`tcp`** ■ **`udp`** To block all protocols, do not include this parameter.
--port-range-min	Enter the first port to block. To block all ports, do not include this parameter. To block a single port, enter the same value for the `--port-range-min` and `--port-range-max` parameters.
--port-range-max	Enter the last port to block. To block all ports, do not include this parameter. To block a single port, enter the same value for the `--port-range-min` and `--port-range-max` parameters.
--remote-ip-prefix	Enter the source network of traffic to block (for example, 10.10.0.0/24). This parameter cannot be used together with the `--remote-group-id` parameter.
--remote-group-id	Enter the name or ID of the source security group of traffic to block. This parameter cannot be used together with the `--remote-ip-prefix` parameter.

The provider security group rules are enforced on all newly created ports on virtual machines in the specified project and cannot be overridden by tenant-defined security groups.

What to do next

You can enforce one or more provider security groups on existing ports by running the following command:

```
neutron port-update port-id --provider-security-groups list=true group-id1...
```

Use NSX Data Center for vSphere Security Policies in OpenStack

You can enforce NSX Data Center for vSphere security policies through Neutron security groups. This feature can also be used to insert third-party network services.

Provider and standard security groups can both consume NSX Data Center for vSphere security policies. Rule-based provider and standard security groups can also be used together with security policy-based security groups. However, a security group associated with a security policy cannot also contain rules.

Security policies take precedence over all security group rules. If more than one security policy is enforced on a port, the order in which the policies are enforced is determined by NSX Data Center for vSphere. You can change the order in the vSphere Web Client on the **Security > Firewall** page under **Networking and Security**.

Prerequisites

Create the desired security policies in NSX Data Center for vSphere. See "Create a Security Policy" in the *NSX Administration Guide*.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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. Uncomment the `nsxv_use_nsx_policies`, `nsxv_default_policy_id`, and `nsxv_allow_tenant_rules_with_policy` parameters and configure them.

Option	Description
nsxv_use_nsx_policies	Enter **`true`**.
nsxv_default_policy_id	Enter the ID of the NSX Data Center for vSphere security policy that you want to associate with the default security group for new projects. If you do not want to use a security policy by default, you can leave this parameter commented out. To find the ID of a security policy, select **Home > Networking & Security** and click **Service Composer**. Open the **Security Policies** tab and click the **Show Columns** icon at the bottom left of the table. Select **Object Id** and click **OK**. The ID of each security policy is displayed in the table.
nsxv_allow_tenant_rules_with_policy	Enter **`true`** to allow tenants to create security groups and rules or **`false`** to prevent tenants from creating security groups or rules.

5. Deploy the updated configuration.

   ```
   sudo viocli deployment configure
   ```

Deploying the configuration briefly interrupts OpenStack services.

6 Log in to the controller node as `viouser`.

7 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

8 If you want to use additional security groups with security policies, you can perform the following steps:

- To associate an NSX Data Center for vSphere security policy with a new security group, create the group and update it with the desired policy:

```
neutron security-group-create security-group-name --tenant-id tenant-uuid
neutron security-group-update --policy=policy-id security-group-uuid
```

- To migrate an existing security group to a security policy-based group, run the following command:

```
sudo -u neutron nsxadmin -r security-groups -o migrate-to-policy --property policy-id=policy-id --property security-group-id=security-group-uuid
```

Note This command removes all rules from the specified security group. Ensure that the target policy is configured such that the network connection will not be interrupted.

9 Configure Neutron to prioritize NSX Data Center for vSphere security policies over security groups.

```
sudo -u neutron nsxadmin --config-file /etc/neutron/neutron.conf --config-file /etc/neutron/plugins/vmware/nsxv.ini -r firewall-sections -o nsx-reorder
```

Authentication and Identity

4

In VMware Integrated OpenStack, authentication and identity management are provided by the Keystone component. In addition to SQL-backed OpenStack users, you can also configure authentication through LDAP or through identity federation with VMware Identity Manager.

This chapter includes the following topics:

Domain Management

You can create domains to manage users and tenants.

All VMware Integrated OpenStack deployments contain the `local` and `Default` domains.

- The `local` domain includes service users and is backed by a local SQL database.
- The `Default` domain contains standard OpenStack users. If you configure LDAP during VMware Integrated OpenStack installation (single domain), the `Default` domain is backed by LDAP and also contains LDAP users. Otherwise, the `Default` is backed by the local SQL database.
- The `admin` user is a member of both `local` and `Default` domains.

Important Do not disable or delete the `local` or `Default` domains.

You can create and manage additional domains as needed. For example, you can create a separate domain for federated users. To manage domains, select **Identity > Domains** on the VMware Integrated OpenStack dashboard.

Configure LDAP Authentication

You can configure LDAP authentication or modify your existing LDAP configuration.

VMware Integrated OpenStack supports SQL plus one or more domains as an identity source, up to a maximum of 10 domains.

Important All LDAP attributes must use ASCII characters only.

Prerequisites

Contact your LDAP administrator or use tools such as ldapsearch or Apache Directory Studio to obtain the correct values for LDAP settings.

Procedure

1. In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.
2. Open the **Manage** tab.
3. On the **Settings** tab, click **Configure Identity Source**.
4. Click the **Add** (plus sign) icon to configure a new LDAP source or the **Edit** (pencil) icon to modify an existing configuration.
5. Enter your LDAP configuration.

Option	Description
Active Directory domain name	Specify the full Active Directory domain name.
Keystone domain name	Enter the Keystone domain name for the LDAP source. Do not use `default` or `local` as the Keystone domain.
Bind user	Enter the user name to bind to Active Directory for LDAP requests.
Bind password	Enter the password to allow the LDAP client access to the LDAP server.
Domain controllers	(Optional) Enter the IP addresses of one or more domain controllers, separated with commas (,). If you do not specify a domain controller, VMware Integrated OpenStack will automatically choose an existing Active Directory domain controller.
Site	(Optional) Enter a specific deployment site within your organization to limit LDAP searching to that site.
User Tree DN	(Optional) Enter the search base for users (for example, **`DC=vmware, DC=com`**). In most Active Directory deployments, the top of the user tree is used by default.
User Filter	(Optional) Enter an LDAP search filter for users. **Important** If 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 more information about filters, see "Search Filter Syntax" in the Microsoft documentation at https://docs.microsoft.com/en-us/windows/win32/adsi/search-filter-syntax.
Group tree DN	(Optional) Enter the search base for groups. The LDAP suffix is used by default.

Option	Description
Group filter	(Optional) Enter an LDAP search filter for groups.
LDAP admin user	Enter an LDAP user to act as an administrator for the domain. If you specify an LDAP admin user, the `admin` project will be created in the Keystone domain for LDAP, and this user will be assigned the `admin` role in that project. This user can then log in to Horizon and perform other operations in the Keystone domain for LDAP. If you do not specify an LDAP admin user, you must use the OpenStack command-line interface to add a project to the Keystone domain for LDAP and assign the `admin` role to an LDAP user in that project.

You can select the **Advanced settings** check box to display additional LDAP configuration fields.

Option	Description
Encryption	Select **None**, **SSL**, or **StartTLS**.
Hostname	Enter the hostname of the LDAP server.
Port	Enter the port number to use 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. 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 the 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.

6 Click **Validate** to confirm your settings.

7 Click **OK**.

8 If you did not specify an LDAP admin user, configure a project and administrator for the Keystone domain for LDAP.

   a Log in to the OpenStack Management Server as `viouser`.

   b Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su –
source ~/cloudadmin.rc
```

   c Create a project in the Keystone domain for LDAP.

```
openstack project create new-project --domain ldap-domain
```

d Add an LDAP user to the new project.

```
openstack user set ldap-username --domain ldap-domain --project new-project --project-domain
ldap-domain
```

e In the Keystone domain for LDAP, assign the `admin` role to the LDAP user.

```
openstack role add admin --user ldap-username --user-domain ldap-domain --domain ldap-domain
```

f In the new project, assign the `admin` role to the LDAP user.

```
openstack role add admin --user ldap-username --user-domain ldap-domain --project new-project
--project-domain ldap-domain
```

LDAP authentication is configured on your VMware Integrated OpenStack deployment. You can log in to the VMware Integrated OpenStack dashboard as the LDAP admin user that you specified during configuration.

Configure VMware Identity Manager Federation

You can configure VMware Integrated OpenStack to use VMware Identity Manager as an identity provider solution.

Users can authenticate with VMware Identity Manager over the Security Association Markup Language (SAML) 2.0 protocol. Federated users must authenticate using the VMware Integrated OpenStack dashboard. The OpenStack command-line interface is not supported.

Prerequisites

- Deploy and configure VMware Identity Manager 2.8 or later.
- Ensure that your VMware Identity Manager instance can communicate with the VMware Integrated OpenStack management network.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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 Add the following parameters.

Option	Description
federation_protocol	Enter **saml2**.
federation_idp_id	Enter a name for the identity provider. This name is used in OpenStack Management Server command-line operations and cannot include special characters or spaces.
federation_idp_name	Enter a display name for the identity provider. This name is shown to users under **Authenticate using** when they log in to the VMware Integrated OpenStack dashboard.
federation_idp_metadata_url	Enter **https://*identity-mgr-fqdn*/SAAS/API/1.0/GET/metadata/idp.xml**.
federation_group	Enter a group to contain federated users.
federation_group_description	Enter a description for the federated users group.
vidm_address	Enter the FQDN of your VMware Identity Manager instance (for example, `https://vxlan-vm-2-10.network.example.com`).
vidm_user	Enter the user name of a VMware Identity Manager administrator.
vidm_password	Enter the password for the VMware Identity Manager administrator.
vidm_insecure	Enter **false** to verify TLS certificates or **true** to disable certificate verification.
vidm_group	Enter the user group in VMware Identity Manager to use for federation.

5 Deploy the updated configuration.

```
sudo viocli deployment configure
```

Deploying the configuration briefly interrupts OpenStack services.

6 Assign projects and roles to federated users or groups.

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 **Identity > Projects**.

d Click **Manage Members** next to the desired project.

e Add federated users or groups and specify the desired roles.

f Click **Save**.

VMware Integrated OpenStack is integrated with VMware Identity Manager, and federated users and groups are imported into OpenStack. When you access the VMware Integrated OpenStack dashboard, you can choose the VMware Identity Manager identity provider to log in as a federated user.

OpenStack Projects and Users 5

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 control network security on the project level through provider security groups or NSX Data Center for vSphere security policies.

This chapter includes the following topics:

Create an OpenStack Project

Projects are organizational units in OpenStack. They can contain users, instances, and other objects such as images.

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 **Identity > Projects** and click **Create Project**.
4. On the **Project Information** tab, enter a name and description and select whether to enable the project.
5. (Optional) On the **Project Members** tab, add users to the project.
6. (Optional) On the **Project Groups** tab, add user groups to the project.
7. On the **Quotas** tab, specify resource quotas for the project.
8. Click **Create Project**.

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

Note The project ID generated is 32 characters in length. However, when filtering by project ID specific to the security group section in Neutron server logs or in vRealize Log Insight, use only the first 22 characters.

What to do next

In the **Actions** column to the right of each project, you can modify project settings, including adding and removing users and groups, modifying project quotas, and changing the name or enabled status of the project.

If you disable a project, it is no longer accessible to its members, but its instances continue to run, and project data is retained. Users that are assigned only to disabled projects cannot log in to the VMware Integrated OpenStack dashboard.

You can select one or more projects and click **Delete Projects** to remove them permanently. Deleted projects cannot be restored.

Create a Cloud User

Cloud users have fewer permissions than cloud administrators. Cloud users can create and manage instances, volumes, networks, and images for the project to which they are assigned.

Prerequisites

Create and enable at least one OpenStack project. See Create an OpenStack Project.

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 **Identity > Users** and click **Create User**.
4. Configure the user settings.

Option	Description
User Name	Enter the user name.
Description	(Optional) Enter a description for the user.
Email	(Optional) Enter an email address for the user.
Password/Confirm Password	Enter a preliminary password for the user. The password can be changed after the user logs in for the first time.
Primary Project	Select the project to which the user is assigned. A user account must be assigned to at least one project.
Role	Select a role for the user. The user inherits the permissions assigned to the specified role.
Enable	Select **Enable** to allow to user to log in and perform OpenStack operations.

5. Click **Create User**.

What to do next

In the **Actions** column to the right of each user, you can modify user settings, change the user password, and enable or disable the user.

If you want to assign a single user to multiple projects, select **Identity > Projects** and click **Manage Members** next to the desired project.

You can create a group containing multiple users for simpler administration. See Create a User Group.

You can select one or more users and click **Delete Users** to remove them permanently. Deleted users cannot be restored.

Create a User Group

You can create a group containing multiple users for easier administration.

Prerequisites

Create the desired users. See Create a Cloud User.

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 **Identity > Groups** and click **Create Group**.
4. Enter a name and description and click **Create Group**.
5. In the **Actions** column to the right of the new group, click **Manage Members**.
6. Click **Add Users**.
7. Select one or more users and click **Add Users**.

What to do next

You can add the user group when you create or modify a project. All users in the group inherit the roles specified in the project for the group.

OpenStack Instances

6

Instances are virtual machines that run in the cloud.

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.

This chapter includes the following topics:

Import Virtual Machines into VMware Integrated OpenStack

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

Imported virtual machines become OpenStack instances but remain distinct.

- If a virtual machine has multiple disks, the disks are imported as Cinder volumes.
- Existing networks are imported as provider networks of type `portgroup` with access restricted to the given tenant.
- After a virtual machine with a specific network backing is imported, the same network cannot be imported to a different project.
- Neutron subnets are automatically created with DHCP disabled.

- Neutron ports are automatically created based on the IP and MAC address of the network interface card on the virtual machine.

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. To avoid this problem, use static DHCP bindings on existing DHCP servers and do not run new OpenStack instances on imported networks.

You import VMs using the Data Center Command-Line Interface (DCLI) on the OpenStack Management Server.

Prerequisites

- Deploy VMware Integrated OpenStack with NSX Data Center for vSphere or VDS networking. Importing virtual machines is not supported for NSX-T Data Center deployments.
- Verify that the virtual machines that you want to import are in the same vCenter Server instance.

Procedure

1. In vSphere, add the clusters containing the desired virtual machines as compute clusters in your VMware Integrated OpenStack deployment. For instructions, see Add Compute Clusters to Your Deployment.
2. Log in to the OpenStack Management Server as `viouser`.
3. If you want to prevent imported virtual machines from being relocated or renamed, update your deployment configuration.
   a. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

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

   b. Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
   c. Uncomment the `nova_import_vm_relocate` parameter and set its value to `false`.
   d. Deploy the updated configuration.

      ```
      sudo viocli deployment configure
      ```

      Deploying the configuration briefly interrupts OpenStack services.
4. Connect to the VMware Integrated OpenStack vAPI endpoint.

   ```
   dcli +server http://mgmt-server-ip:9449/api +i
   ```

   If you cannot connect to the server, see DCLI Cannot Connect to Server.

5 Import unmanaged virtual machines into VMware Integrated OpenStack.

Note When you execute a command, DCLI prompts you to enter the administrator credentials for your vCenter Server instance. You can save these credentials to avoid entering your username and password every time.

- Run the following command to import all unmanaged virtual machines:

```
com vmware vio vm unmanaged importall --cluster cluster-name [--tenant-mapping {FOLDER |
RESOURCE_POOL} [--root-folder root-folder | --root-resource-pool root-resource-pool]]
```

Option	Description
`--cluster`	Enter the compute cluster that contains the virtual machines that you want to import.
`--tenant-mapping {FOLDER | RESOURCE_POOL}`	Specify whether to map imported virtual machines to OpenStack projects based on their location in folders or resource pools. If you do not include this parameter, all imported VMs will become instances in the **import_service** project by default.
`--root-folder ROOT_FOLDER`	If you specified **FOLDER** for the `--tenant-mapping` parameter, you can provide the name of the root folder containing the virtual machines to be imported. All virtual machines in the specified folder or any of its subfolders are imported as instances into an OpenStack project with the same name as the folder in which they are located. **Note** If you specify `--tenant-mapping FOLDER` but do not specify `--root-folder`, the name of the top-level folder in the cluster is used by default.
`--root-resource-pool ROOT_RESOURCE_POOL`	If you specified **RESOURCE_POOL** for the `--tenant-mapping` parameter, you can provide the name of the root resource pool containing the virtual machines to be imported. All virtual machines in the specified resource pool or any of its child resource pools are imported as instances into an OpenStack project with the same name as the resource pool in which they are located.

- Run the following command to import a specified virtual machine:

```
com vmware vio vm unmanaged importvm --vm vm-id [--tenant project-name] [--nic-mac-address
nic-mac --nic-ipv4-address nic-ip] [--root-disk root-disk-path] [--nics specifications]
```

Option	Description
`--vm`	Enter the identifier of the virtual machine that you want to import. You can view the ID values of all unmanaged virtual machines by running the `com vmware vio vm unmanaged list` command.
`--tenant`	Specify the OpenStack project into which you want to import the virtual machine. If you do not include this parameter, the `import_service` project is used by default.
`--nic-mac-address`	Enter the MAC address of the network interface card on the virtual machine. If you do not include this parameter, the import process attempts to discover the MAC and IP addresses automatically. **Note** If you include this parameter, you must also include the `nic_ipv4_address` parameter.

Option	Description
`--nic-ipv4-address`	Enter the IP address and prefix for the network interface card on the virtual machine. Enter the value in CIDR notation (for example, 10.10.1.1/24). This parameter must be used together with the `--nic-mac-address` parameter.
`--root-disk`	For a virtual machine with multiple disks, specify the root disk datastore path in the following format: **`--root-disk '[datastore1] foo/foo_1.vmdk'`**
`--nics`	For a virtual machine with multiple NICs, specify the MAC and IP addresses of each NIC in JSON format. Use the following key-value pairs: ■ mac_address: MAC address of the NIC in standard format ■ ipv4_address: IPv4 address in CIDR notation For example: `--nics '[{"mac_address": "00:50:56:9a:f5:7b", "ipv4_address": "10.10.1.1/24"}, {"mac_address": "00:50:56:9a:ee:be", "ipv4_address": "10.10.2.1/24"}]'`

Migrate an Instance

You can live-migrate an OpenStack instance to a different compute node.

Note Instances managed by VMware Integrated OpenStack must be migrated by using OpenStack commands. Do not use vCenter Server or other methods to migrate OpenStack instances.

Prerequisites

- The source and target compute nodes must both be located within the same vCenter Server instance.
- The source and target compute nodes must have at least one distributed switch in common. If two distributed switches are attached to the source compute node but only one distributed switch is attached to the target compute node, live migration will succeed but the OpenStack instance will be connected only to the port group of the distributed switch common to both compute nodes.
- Instances with a CD-ROM drive attached cannot be live-migrated.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Log in to the controller node as `viouser`.
3. Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

4. If a volume is attached to the instance, prepare the volume for migration.

```
viocli ds-migrate-prep datacenter-name datastore-name
```

Specify the name of the vSphere data center and datastore that contain the volume attached to the instance.

5 Migrate the instance to the desired compute node.

```
openstack server migrate compute-name instance-uuid --live
```

- To find the name of a compute node, run the `openstack host list` command and view the **Host Name** column.
- To find the UUID of the instance, run the `openstack server list` command and view the **ID** column.

What to do next

You can run the `openstack server show instance-uuid` command to confirm that the instance has been migrated to the desired compute node.

Enable Live Resize

You can enable live resize for OpenStack instances by configuring image metadata. With live resize, you can change the disk size, memory, and vCPUs of an instance while the instance is powered on.

Prerequisites

- Do not create live resize-enabled instances using SR-IOV-enabled ports. Live resize is not compatible with SR-IOV.
- Do not use live resize-enabled instances in tenant virtual data centers. Live resize is not compatible with tenant virtual data centers.

In addition, the following conditions apply for live resizing of disk size:

- Use VMDK as the disk format for the image.
- Use a SCSI virtual disk adapter type for the image. IDE adapter types are not supported.
- Deploy virtual machines from the image as full clones. Linked clones cannot be live-resized.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/clouddadmin.rc
```

3 Create a new image with live resize enabled.

```
openstack image create image-name --disk-format {vmdk | iso} --container-format bare --file image-file {--public | --private} [--property vmware_adaptertype="vmdk-adapter-type"] [--property vmware_disktype="{sparse | preallocated | streamOptimized}"] --property vmware_ostype="operating-system" --property img_linked_clone="false" --property os_live_resize="{vcpu | memory | disk}"
```

Option	Description
image-name	Enter the name of the source image.
--disk-format	Enter **vmdk**.
--container-format	Enter **bare**. The container format argument is not currently used by Glance.
--file	Specify the image file to upload.
{--public | --private}	Include `--public` to make the image available to all users or `--private` to make the image available only to the current user.
--property vmware_adaptertype	Specify the adapter type of the VMDK disk. For disk live resize, you must specify a SCSI adapter. If you do not include this parameter, the adapter type is determined by introspection.
--property vmware_disktype	Specify **sparse**, **preallocated**, or **streamOptimized**. If you do not include this parameter, the disk type is determined by introspection.
--property vmware_ostype	Specify the operating system on the image.
--property img_linked_clone	Enter **false**.
--property os_live_resize	Specify **vcpu**, **memory**, **disk**, or any combination separated by commas (for example, **vcpu,memory,disk**).

When you create virtual machines using the image that you defined in this procedure, those virtual machines can be resized without needing to be powered off.

Use Affinity to Control OpenStack Instance Placement

You can place instances using OpenStack server groups with an affinity or anti-affinity policy. Affinity indicates that all instances in the group must placed on the same host, and anti-affinity indicates that no instances in the group can be placed on the same host.

Affinity and anti-affinity policies cannot determine the specific ESXi host on which instances are placed. These policies only control whether instances are placed on the same hosts as other instances in a server group. To place instances on specific hosts, see Use DRS to Control OpenStack Instance Placement.

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.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

3 Create a server group with the desired policy.

```
openstack server group create group-name --policy {affinity | anti-affinity}
```

Option	Description
group-name	Enter a name for the server group.
--policy	Enter **affinity** to place instances on the same host or **anti-affinity** to prevent instances from being placed on the same host.

4 When you launch an instance, pass the server group as a scheduler hint to implement affinity or anti-affinity.

```
openstack server create instance-name --image image-uuid --flavor flavor-name --nic net-
id=network-uuid --hint group=servergroup-uuid
```

What to do next

Confirm that the affinity rules and instances are configured correctly. In vCenter Server, select the compute cluster, open the **Configure** tab, and click **VM/Host Rules**.

Use DRS to Control OpenStack Instance Placement

You can use vSphere DRS settings to control how OpenStack instances are placed on hosts in the compute cluster. After configuring DRS, you modify the metadata of source images in OpenStack to ensure that instances generated from those images are correctly identified for placement.

Procedure

1 Define VM and Host Groups for Placing OpenStack Instances

You define VM and host groups to contain and manage specific OpenStack instances.

2 Create a DRS Rule for OpenStack Instance Placement

You 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 modify the metadata of a source image to automatically place instances into VM groups. DRS rules then determine the host groups on which these instances will be created.

Define VM and Host Groups for Placing OpenStack Instances

You define VM and host groups to contain and manage specific OpenStack instances.

Prerequisites

- Ensure that the compute cluster contains at least one virtual machine. If the compute cluster does not contain any virtual machines, create a dummy virtual machine for this procedure.
- On the compute cluster, enable DRS and set **DRS Automation** to **Partially automated** or **Fully automated**.
- On the compute cluster, set **Power Management** to **Off**.

Procedure

1. In the vSphere Web Client, select the compute cluster and click **Configure**.
2. Under **Configuration**, click **VM/Host Groups**.
3. Create a VM group.
   a. Click **Add**.
   b. Enter a name and select **VM Group** from the **Type** drop-down menu.
   c. Click **Add**.
   d. On the **Filter** tab, select virtual machines to add to the group.
   e. Click **OK**.
4. Create a host group.
   a. Click **Add**.
   b. Enter a name and select **Host Group** from the **Type** drop-down menu.
   c. Click **Add**.
   d. On the **Filter** tab, select hosts to add to the group.
   e. Click **OK**.

What to do next

Create a rule that determines how OpenStack instances assigned to the VM group are distributed on the hosts in the host group.

Create a DRS Rule for OpenStack Instance Placement

You create a DRS rule to manage the distribution of OpenStack instances in a VM group to a specific host group.

Prerequisites

- Define at least one VM group and at least one host group. See Define VM and Host Groups for Placing OpenStack Instances.

- On the compute cluster, enable DRS and set **DRS Automation** to **Partially automated** or **Fully automated**.
- On the compute cluster, set **Power Management** to **Off**.

Procedure

1. In the vSphere Web Client, click the compute cluster and select **Configure**.
2. Under **Configuration**, click **VM/Host Rules**.
3. Click the **Add...** button.
4. Enter a name for the rule and select the **Enable rule** option.
5. In the **Type** drop-down menu, select **Virtual Machines to Hosts**.
6. In the **VM Group** drop-down menu, select the VM group that contains the OpenStack instances you want to place.
7. In the next drop-down menu, select a specification for the rule.

Option	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 hosts 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.

8. In the **Host Group** drop-down menu, select the host group that contains the hosts on which the OpenStack instances will be placed and click **OK**.

What to do next

In the VMware Integrated OpenStack dashboard, you can 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 modify the metadata of a source image to automatically place instances into VM groups. DRS rules then determine the host groups on which these instances will be created.

Prerequisites

Configure a VM group and host group for the compute cluster.

Procedure

1. Log in to the VMware Integrated OpenStack dashboard as a cloud administrator and select the `admin` project from the drop-down menu in the title bar.
2. Select **Admin > System > Images**.
3. Create a new image or choose an existing image.
4. Click the down arrow next to the flavor that you want to use and select **Update Metadata**.
5. In the **Available Metadata** pane, expand **VMware Driver Options** and click the **Add** (plus sign) icon next to **DRS VM group**.
6. Enter the desired VM group name as the value of the `vmware_vm_group` parameter and click **Save**.

All OpenStack instances generated from this source image will be automatically assigned to the specified VM group and governed by its DRS rules.

Configure QoS Resource Allocation for Instances

You can control resource allocations for CPU, memory, disk IOPS, and virtual network interfaces by modifying a flavor or image.

Note Configuring virtual interface quotas is not supported in NSX-T Data Center. The VIF limit, reservation, and shares settings cannot be used with NSX-T Data Center deployments.

QoS resource allocation can also be specified by flavor extra specs or image metadata. If flavor and image settings conflict, the image metadata configuration takes precedence.

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. Specify a flavor or image to use for QoS.
   - To use flavor extra specs for QoS configuration, perform the following steps:
     a. Select **Admin > System > Flavors**.
     b. Create a new flavor or choose an existing flavor to use for QoS.
     c. Select **Update Metadata** next to the flavor that you want to use.
   - To use image metadata for QoS configuration, perform the following steps:
     a. Select **Admin > System > Images**.
     b. Create a new image or choose an existing image to use for QoS.
     c. Click the down arrow next to the image that you want to use and select **Update Metadata**.
4. In the **Available Metadata** pane, expand **VMware Quota**.

5 Click the **Add** (plus sign) icon next to the item that you want to use.

Metadata Property	Description
Quota: CPU Limit	Specify the maximum CPU allocation in megahertz. The value `0` indicates that CPU usage is not limited.
Quota: CPU Reservation	Specify the guaranteed CPU allocation in megahertz.
Quota: CPU Shares Level	Specify the level of CPU shares allocated. You can enter **`custom`** and add the CPU Shares Value quota to provide a custom value.
Quota: CPU Shares Value	Specify the number of CPU shares allocated. If the CPU Shares Level quota is not set to `custom`, this value is ignored.
Quota: Disk IO Limit	Specify the maximum disk transaction allocation in IOPS. The value `0` indicates that disk transactions are not limited.
Quota: Disk IO Reservation	Specify the guaranteed disk transaction allocation in IOPS.
Quota: Disk IO Shares Level	Specify the level of disk transaction shares allocated. You can enter **`custom`** and add the Disk IO Shares Value quota to provide a custom value.
Quota: Disk IO Shares Value	Specify the number of disk transaction shares allocated. If the Disk IO Shares Level quota is not set to `custom`, this value is ignored.
Quota: Memory Limit	Specify the maximum memory allocation in megabytes. The value `0` indicates that memory usage is not limited.
Quota: Memory Reservation	Specify the guaranteed memory allocation in megabytes.
Quota: Memory Shares Level	Specify the level of memory shares allocated. You can enter **`custom`** and add the Memory Shares Value quota to provide a custom value.
Quota: Memory Shares Value	Specify the number of memory shares allocated. If the Memory Shares Level quota is not set to `custom`, this value is ignored.
Quota: VIF Limit	Specify the maximum virtual interface bandwidth allocation in Mbps. The value `0` indicates that virtual interface bandwidth is not limited.
Quota: VIF Reservation	Specify the guaranteed virtual interface bandwidth allocation in Mbps.
Quota: VIF Shares Level	Specify the level of virtual interface bandwidth shares allocated. You can enter **`custom`** and add the VIF Shares Value quota to provide a custom value.
Quota: VIF Shares Value	Specify the number of virtual interface bandwidth shares allocated. If the VIF Shares Level quota is not set to `custom`, this value is ignored.

6 Click **Save**.

You can now deploy QoS-enabled instances by configuring them with the flavor or image that you modified in this procedure.

To apply QoS settings to an existing instance, resize the instance and select the flavor with the desired QoS settings. The specified settings take effect after the resize process is complete.

Use Storage Policy-Based Management with OpenStack Instances

You can use vSphere storage policies to control the datastores on which OpenStack instances are created.

Prerequisites

Create the desired storage policy in vSphere.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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. Uncomment the `nova_pbm_enabled` parameter and set its value to **`true`**.
5. Uncomment the `nova_pbm_default_policy` parameter and set its value to the name of the storage policy to use by default when an instance is created with a flavor that is not associated with a storage policy.
6. Uncomment the `nova_scheduler_default_filters` parameter and add **`AggregateInstanceExtraSpecsFilter`** to the end.

   ```
   nova_scheduler_default_filters: RetryFilter, AvailabilityZoneFilter, RamFilter, ComputeFilter,
   ComputeCapabilitiesFilter, ImagePropertiesFilter, ServerGroupAntiAffinityFilter,
   ServerGroupAffinityFilter, PciPassthroughFilter, AggregateInstanceExtraSpecsFilter
   ```

7. Uncomment the `nova_use_linked_clone` parameter and set its value to **`false`**.
8. Deploy the updated configuration.

   ```
   sudo viocli deployment configure
   ```

   Deploying the configuration briefly interrupts OpenStack services.
9. Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
10. Select the **admin** project from the drop-down menu in the title bar.
11. Select **Admin > System > Flavors**.
12. Create a new flavor or choose an existing flavor.
13. Click **Update Metadata** to the right of the flavor.

14 In the **Available Metadata** pane, expand **VMware Policies** and click the **Add** (plus sign) icon next to **Storage Policy**.

15 Enter the desired storage policy name as the value of the `vmware:storage_policy` parameter and click **Save**.

The specified vSphere storage policy is applied to all new OpenStack instances that are created from the flavor. The default storage policy is applied to all new instances that are created from a flavor not associated with a storage policy.

Configure Virtual CPU Pinning

When running latency-sensitive applications inside a virtual machine, you can use virtual CPU pinning to eliminate the extra latency that is imposed by virtualization.

Important This feature is offered in VMware Integrated OpenStack Carrier Edition only. For more information, see "VMware Integrated OpenStack Licensing" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

Virtual CPU pinning enables high latency sensitivity and ensures that all memory and an entire physical core are reserved for the virtual CPU of an OpenStack instance. You configure virtual CPU pinning on a flavor and then create instances with that flavor.

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 Create a new flavor or choose an existing flavor to use for virtual CPU pinning.

5 Select **Update Metadata** next to the flavor that you want to use.

6 In the **Available Metadata** pane, select and configure the required metadata.

   a Expand **CPU Pinning** and click the **Add** (plus sign) icon next to **CPU Pinning policy**.

   b Set the value of `hw:cpu_policy` to **`dedicated`**.

   c Expand **VMware Policies** and click the **Add** (plus sign) icon next to **VM latency sensitivity**.

   d Set the value of `vmware:latency_sensitivity_level` to **`high`**.

   e Expand **VMware Quota** and click the **Add** (plus sign) icon next to **CPU Reservation in Percentage** and **Memory Reservation in Percentage**.

   f Set the value of `quota:cpu_reservation_percent` and `quota:memory_reservation_percent` to **`100`**.

7 Click **Save**.

What to do next

You can now enable virtual CPU pinning on an instance by configuring it with the flavor that you modified in this procedure.

Configure OpenStack Instances for NUMA

VMware Integrated OpenStack supports non-uniform memory access (NUMA)-aware placement of OpenStack instances on the underlying vSphere environment.

Important This feature is offered in VMware Integrated OpenStack Carrier Edition only. For more information, see "VMware Integrated OpenStack Licensing" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

NUMA links small, cost-effective nodes using a high-performance connection to provide low latency and high throughput. This performance is often required for virtual network functions (VNFs) in telecommunications environments. For information about NUMA in vSphere, see "Using NUMA Instances with ESXi" in *vSphere Resource Management*.

To obtain information about your current NUMA configuration, run the following command on your ESXi hosts:

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

Prerequisites

- Ensure that vCPUs, memory, and physical NICs intended for virtual machine traffic are placed on same node.
- In vSphere, create a teaming policy that includes all physical NICs on the NUMA node. See "Teaming and Failover Policy" in *vSphere Networking*.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Switch to the `root` user and load the cloud administrator credentials file.

   ```
   sudo su –
   source ~/cloudadmin.rc
   ```

3. Create a Neutron network on which all physical NICs are located on a single NUMA node.
4. Create an OpenStack flavor that includes the `numa.nodeAffinity` property.

   ```
   nova flavor-key flavor-id set vmware:extra_config='{"numa.nodeAffinity": "numa-node-id"}'
   ```

5. Launch an OpenStack instance using the flavor and network created in this procedure.

Configuring Passthrough Devices on OpenStack Instances

You can create OpenStack instances using DirectPath I/O and Single Root I/O Virtualization (SR-IOV) passthrough devices.

Important This feature is offered in VMware Integrated OpenStack Carrier Edition only. For more information, see "VMware Integrated OpenStack Licensing" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

Passthrough associates a physical device with a virtual machine, reducing the latency caused by virtualization. The following table shows how passthrough is implemented in VMware Integrated OpenStack.

Table 6-1. Key Passthrough Components and Roles

Component	Role
Nova compute	■ 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	■ 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	■ Schedules instance placement on hosts that match the PCI requests
vSphere	■ Manages hosts in a dedicated compute cluster with NICs and hosts enabled for SR-IOV. **Note** DRS rules do not apply to devices enabled for SR-IOV. Place SR-IOV hosts in a separate compute cluster.

Configure Passthrough for Networking Devices

You can configure a port to allow SR-IOV or DirectPath I/O passthrough and then create OpenStack instances that use physical hardware interfaces.

Important This feature is offered in VMware Integrated OpenStack Carrier Edition only. For more information, see "VMware Integrated OpenStack Licensing" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

This procedure uses OpenStack Neutron to enable passthrough for networking devices. For non-networking devices, see Configure Passthrough for Non-Networking Devices.

Prerequisites

- Verify that your OpenStack deployment is using VDS or NSX Data Center for vSphere networking. Deployments with NSX-T Data Center do not support passthrough.

- Enable SR-IOV or DirectPath I/O in vSphere:
  - To enable SR-IOV, see "Enable SR-IOV on a Host Physical Adapter" in *vSphere Networking*.
  - To enable DirectPath I/O, see "Enable Passthrough for a Network Device on a Host" in *vSphere Networking*.
- Create a dedicated compute cluster for SR-IOV devices. DRS rules do not apply to these devices.
- To persist the MAC address of a physical device, add its cluster as a compute node before enabling direct passthrough on the device. If direct passthrough has already been enabled, you can disable it, restart the cluster, and enable direct passthrough again.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

3. Log in to the OpenStack Management Server.
4. Create a provider network for SR-IOV devices.

```
neutron net-create network-name --tenant-id project-uuid --provider:network_type {vlan |
portgroup | nsx-net} --provider:physical_network physical-id [--provider:segmentation_id vlan-id]
```

Option	Description
network-name	Enter a name for the network.
--tenant-id	Specify the UUID of the project for which to create the port. You can find the UUID of a project by running the `openstack project list` command.
--provider:network_type	Enter **`vlan`** or **`portgroup`**.
--provider:physical_network	■ For a VLAN network, specify the managed object identifier (MOID) of the distributed switch. ■ For a port group network, specify the MOID of the port group.
--provider:segmentation_id	If you want to create a VLAN-based network, enter the VLAN ID.

5. Create a passthrough-enabled port.

```
neutron port-create network-id --tenant-id project-uuid --name port-name --vnic_type {direct |
direct-physical}
```

Option	Description
network-id	Specify the UUID of the network on which to create the port. You can find the UUID of a network by running the `openstack network list` command.
--tenant-id	Specify the UUID of the project for which to create the port. You can find the UUID of a project by running the `openstack project list` command.

Option	Description
--name	Enter a name for the port.
--vnic_type	Enter **`direct`** for SR-IOV or **`direct-physical`** for direct passthrough.

Note Port security is not supported for `direct` and `direct-physical` ports and will be automatically disabled for the port created.

You can now deploy passthrough-enabled virtual machines by configuring them with the port that you created during this procedure.

Configure Passthrough for Non-Networking Devices

You can configure flavor and image metadata to allow SR-IOV or DirectPath I/O passthrough and then create OpenStack instances that use physical hardware interfaces.

Important This feature is offered in VMware Integrated OpenStack Carrier Edition only. For more information, see "VMware Integrated OpenStack Licensing" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

This procedure uses OpenStack Nova to enable passthrough for non-networking devices. For networking devices, see Configure Passthrough for Networking Devices .

Prerequisites

- Verify that your OpenStack deployment is using VDS or NSX Data Center for vSphere networking. Deployments with NSX-T Data Center do not support passthrough.
- Enable SR-IOV or DirectPath I/O in vSphere:
  - To enable SR-IOV, see "Enable SR-IOV on a Host Physical Adapter" in *vSphere Networking*.
  - To enable DirectPath I/O, see "Enable Passthrough for a Network Device on a Host" in *vSphere Networking*.
- Create a dedicated compute cluster for SR-IOV devices. DRS rules do not apply to these devices.
- To persist the MAC address of a physical device, add its cluster as a compute node before enabling direct passthrough on the device. If direct passthrough has already been enabled, you can disable it, restart the cluster, and enable direct passthrough again.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

```
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 Uncomment the `nova_pci_alias` parameter and modify its value to match your device.

```
nova_pci_alias: [{"device_type": "type-VF", "name": "virtual-device-name"}, {"vendor_id":"vid",
"product_id":"pid", "device_type": "type-PF", "name":"physical-device-name"}]
```

where:

- `name` (first occurrence) is the alias of the virtual device
- `vendor_id` is the four-digit identifier of the physical device vendor
- `device_id` is the four-digit identifier of the physical device
- `name` (second occurrence) is the alias of the physical device

5 Deploy the updated configuration.

```
sudo viocli deployment configure
```

Deploying the configuration briefly interrupts OpenStack services.

6 Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.

7 Select the **admin** project from the drop-down menu in the title bar.

8 Create a flavor with passthrough enabled.

a Select **Admin > System > Flavors**.

b Create a new flavor or choose an existing flavor to use for passthrough.

c Select **Update Metadata** next to the flavor that you want to use.

d In the **Available Metadata** pane, expand **VMware Driver Options for Flavors** and click the **Add** (plus sign) icon next to **PCI Passthrough alias**.

e Set the value of `pci_passthrough:alias` to *virtual-device-name*:*device-count* and click **Save**.

Option	Description
virtual-device-name	The virtual device name that you specified in Step 4 of this procedure.
device-count	The number of virtual functions that can be called in one request. This value can range from 1 to 10.

9 Create an image with passthrough enabled.

a Select **Admin > System > Images**.

b Create a new image or choose an existing image to use for passthrough.

c Click the down arrow next to the flavor that you want to use and select **Update Metadata**.

d In the **Available Metadata** pane, expand **VMware Driver Options** and click the **Add** (plus sign) icon next to **Virtual Network Interface**.

e Select your device from the drop-down list next to the `hw_vif_model` parameter and click **Save**.

You can now deploy passthrough-enabled virtual machines by configuring them with the flavor and image that you modified during this procedure.

OpenStack Flavors

7

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 chapter includes the following topics:

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.

Supported Flavor Extra Specs

Flavor extra specs are used for advanced configuration of compute instances. VMware Integrated OpenStack exposes additional capabilities through flavor extra specs.

Note Configuring virtual interface quotas is not supported in NSX-T Data Center. The following extra specs cannot be used with NSX-T Data Center deployments:

- `quota:vif_limit`
- `quota:vif_reservation`
- `quota:vif_shares_level`
- `quota:vif_shares_share`

If an image metadata and flavor extra spec conflict, the image metadata takes precedence over the flavor extra spec.

Table 7-1. Flavor Extra Specs in VMware Integrated OpenStack

Extra Spec	Description
`hw:vifs_multi_thread`	Specify **`true`** to provide each virtual interface with its own transmit thread.
`quota:cpu_limit`	Specify the maximum CPU allocation in MHz. The value `0` indicates that CPU usage is not limited.
`quota:cpu_reservation`	Specify the guaranteed CPU allocation in MHz.
`quota:cpu_reservation_percent`	Specify the guaranteed CPU allocation as a percentage of the actual CPU speed of the instance. This parameter takes precedence over the `cpu_reservation` parameter.
`quota:cpu_shares_level`	Specify the level of CPU shares allocated. You can enter **`custom`** and add the `cpu_shares_share` parameter to provide a custom value.
`quota:cpu_shares_share`	Specify the number of CPU shares allocated. If the `cpu_shares_level` parameter is not set to `custom`, this value is ignored.
`quota:disk_io_limit`	Specify the maximum disk transaction allocation in IOPS. The value `0` indicates that disk transactions are not limited.
`quota:disk_io_reservation`	Specify the guaranteed disk transaction allocation in IOPS.

Table 7-1. Flavor Extra Specs in VMware Integrated OpenStack (continued)

Extra Spec	Description
quota:disk_io_shares_level	Specify the level of disk transaction shares allocated. You can enter **custom** and add the disk_io_shares_share parameter to provide a custom value.
quota:disk_io_shares_share	Specify the number of disk transaction shares allocated. If the disk_io_shares_level parameter is not set to custom, this value is ignored.
quota:memory_limit	Specify the maximum memory allocation in MB. The value 0 indicates that memory usage is not limited.
quota:memory_reservation	Specify the guaranteed memory allocation in MB.
quota:memory_reservation_percent	Specify the guaranteed memory allocation as a percentage of the actual memory of the instance. The value 100 indicates that guest memory is also fully reserved. This parameter takes precedence over the memory_reservation parameter.
quota:memory_shares_level	Specify the level of memory shares allocated. You can enter **custom** and add the memory_shares_share parameter to provide a custom value.
quota:memory_shares_share	Specify the number of memory shares allocated. If the memory_shares_level parameter is not set to custom, this value is ignored.
quota:vif_limit	Specify the maximum virtual interface bandwidth allocation in Mbps. The value 0 indicates that virtual interface bandwidth is not limited.
quota:vif_reservation	Specify the guaranteed virtual interface bandwidth allocation in Mbps.
quota:vif_shares_level	Specify the level of virtual interface bandwidth shares allocated. You can enter **custom** and add the vif_shares_share parameter to provide a custom value.
quota:vif_shares_share	Specify the number of virtual interface bandwidth shares allocated. If the vif_shares_level parameter is not set to custom, this value is ignored.
vmware:extra_config	Specify custom configurations in JSON format. For example, '{“acpi.smbiosVersion2.7”:“FALSE”}' .
vmware:hw_version	Specify the hardware version used to create images. In an environment with different host versions, you can use this parameter to place instances on the correct hosts.
vmware:latency_sensitivity_level	Specify the latency sensitivity level for virtual machines.

Table 7-1. Flavor Extra Specs in VMware Integrated OpenStack (continued)

Extra Spec	Description
vmware:resource_pool	Specify the resource pool on which to place new instances. If the name of the project containing the instance matches the name of a resource pool in your environment, the instance is placed in that resource pool by default. Setting this parameter overrides the default behavior and forces the instance to be placed in the specified resource pool.
vmware:storage_policy	Specify the storage policy used for new instances. If Storage Policy-Based Management (SPBM) is not enabled, this parameter is ignored.
vmware:tenant_vdc	Specify the UUID of the tenant virtual data center in which to place instances.
vmware:vm_group	Specify the DRS VM group in which virtual machines will be placed. If the specified VM group does not exist, instances will fail to power on.

Cinder Volumes and Volume Types

8

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

As a cloud administrator, 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 cloud user operations, see "Working with Volumes" in the *VMware Integrated OpenStack User Guide*.

This chapter includes the following topics:

Create a Volume Type

You can create volume types and expose them to one or more tenants for use in volume creation. Volume types can define a vSphere storage profile and default adapter type.

Note Barbican encryption is not supported for volumes or volume types.

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 > Volume Types** and click **Create Volume Type**.
4. Enter a name and description for the volume type.
5. If you want to make the volume type available to certain projects only, deselect **Public**.

   You can configure access to the volume type after it is created.
6. Click **Create Volume Type**

   The new volume type is displayed in the **Volume Types** list.

7 If you want to associate a vSphere storage profile with the volume type, perform the following steps:
   a In the **Actions** column, select **View Extra Specs**.
   b Click **Create**.
   c Enter **`vmware:storage_profile`** in the **Key** text box.
   d Enter the name of the vSphere storage profile in the **Value** text box.
   e Click **Create**.
8 If you want to set a default adapter for the volume type, perform the following steps:
   a In the **Actions** column, select **View Extra Specs**.
   b Click **Create**.
   c Enter **`vmware:adapter_type`** in the **Key** text box.
   d Enter the adapter type in the **Value** text box.

     The following values are supported: **`lsiLogic`**, **`busLogic`**, **`lsiLogicsas`**, **`paraVirtual`**, and **`ide`**.
   e Click **Create**.
9 If your volume type is not public, select **Edit Access** in the **Actions** column and specify the projects that can use the volume type.

   If you do not specify any projects, the volume type is visible only to cloud administrators.

Tenants can select a volume type when creating a volume or modifying an existing volume. The settings defined by the specified volume type are then applied to the new volume.

What to do next

If you want to change the name or description of a volume type, click **Edit Volume Type** in the **Actions** column and make the desired changes. To delete unneeded volume types, select them in the **Volume Types** table and click **Delete Volume Types**.

Modify the Default Volume Adapter Type

By default, empty volumes are created with the LSI Logic Parallel adapter type. You can change the default adapter type for new volumes by modifying your deployment configuration.

Note When a volume is created from an image, the value of the `vmware_adaptertype` metadata in the image is used to determine the adapter type for the volume.

VMware Integrated OpenStack supports the following adapter types:

- LSI Logic Parallel: `lsiLogic`
- BusLogic Parallel: `busLogic`
- LSI Logic SAS: `lsiLogicsas`

- VMware Paravirtual SCSI: `paraVirtual`
- IDE: `ide`

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

```
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. Uncomment the `cinder_volume_default_adapter_type` parameter and set its value to the desired adapter type.
5. Deploy the updated configuration.

```
sudo viocli deployment configure
```

   Deploying the configuration briefly interrupts OpenStack services.

New volumes use the specified adapter type by default.

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.

Note You cannot migrate a volume that has snapshots attached. You must detach all snapshots before migrating a volume.

Migrate All Volumes from a 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.
- On the datastore cluster, enable Storage DRS and set it to **No Automation (Manual Mode)**.
- Detach all snapshots from all volumes on the datastore.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.

2 Prepare the volumes in the datastore for migration.

```
sudo viocli ds-migrate-prep dc-name ds-name
```

Option	Description
dc-name	Enter the name of the data center that contains the desired datastore.
ds-name	Enter the name of the datastore.

3 Place the datastore in maintenance mode.

See "Place a Datastore in Maintenance Mode" in the *vSphere Resource Management* document.

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

Detach all snapshots from the volumes that you want to migrate.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Migrate the volumes.

- To migrate all volumes from a datastore, run the following command:

```
sudo viocli volume-migrate --source-dc src-dc-name --source-ds src-ds-name dest-dc-name dest-ds-name [--ignore-storage-policy]
```

- To migrate specified volumes from a datastore, run the following command:

```
sudo viocli volume-migrate --volume-ids UUID1 dest-dc-name dest-ds-name [--ignore-storage-policy]
```

Option	Description
--source-dc	Enter the data center containing the volumes that you want to migrate. This parameter must be used together with the `--source-ds` parameter. If you want to migrate specified volumes only, do not include this parameter.
--source-ds	Enter the datastore containing the volumes that you want to migrate. This parameter must be used together with the `--source-dc` parameter. If you want to migrate specified volumes only, do not include this parameter.
--volume-ids	Enter the UUID of the volume that you want to migrate. You can include multiple UUIDs separated by commas (,). If you want to migrate all volumes from a datastore, use the `--source-dc` and `--source-ds` parameters instead of this parameter.

Option	Description
dest-dc-name	Enter the name of the data center that contains the datastore to which you want to migrate volumes.
dest-ds-name	Enter the name of the datastore to which you want to migrate volumes.
`--ignore-storage-policy`	Include this parameter to migrate volumes to the target datastore even if the datastore does not comply with the storage policy of the volume.

The specified volumes are migrated to the target datastore.

Migrate Attached Cinder Volumes

You can migrate Cinder volumes that are attached to an OpenStack instance by migrating the corresponding virtual machine to a different datastore.

Note After an attached volume is migrated, the corresponding shadow virtual machine remains on the original datastore but has no disk. When you detach the volume, the disk will be re-attached to the shadow virtual machine.

Prerequisites

Detach all snapshots from the volumes that you want to migrate.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Prepare the datastore containing the volume for migration.

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

```
sudo viocli ds-migrate-prep dc-name ds-name
```

Option	Description
dc-name	Enter the data center that contains the desired volume.
ds-name	Enter the datastore that contains the desired volume.

3 Migrate the instance to which the volume is attached.

```
openstack server migrate compute-name instance-uuid --live
```

- To find the name of a compute node, run the `openstack host list` command and view the **Host Name** column.
- To find the UUID of the instance, run the `openstack server list` command and view the **ID** column.

For more information, see Migrate an Instance.

4 In the vSphere Web Client, migrate the shadow virtual machine for the volume.

   For information, see "Migrate a Virtual Machine to New Storage in the vSphere Web Client" in the *vCenter Server and Host Management* document.

5 If you want to migrate the shadow virtual machine to a cluster in a different availability zone, update the availability zone for the volume.

```
sudo -u cinder cinder-manage volume update_volume_host --volume_id volume-uuid --newhost new-volume-host
```

The Cinder volume and the disk of the corresponding shadow virtual machine are migrated to the new datastore.

Supported Volume Type Extra Specs

Volume type extra specs are used for advanced configuration of Cinder volumes. VMware Integrated OpenStack exposes additional capabilities through volume type extra specs.

Table 8-1. Volume Type Extra Specs in VMware Integrated OpenStack

Extra Spec	Description
`vmware:vmdk_type`	Specify the provisioning format of Cinder volumes in vSphere. You can specify the following formats ■ Thin provision: `thin` ■ Thick provision lazy zeroed: `thick` ■ Thick provision eager zeroed: `eagerZeroedThick`
`vmware:clone_type`	Specify the clone type. You can specify the following types: ■ Full clone: `full` ■ Linked clone: `linked`
`vmware:storage_profile`	Enter the name of the storage policy to use for new volumes.
`vmware:adapter_type`	Specify the adapter type used to attach the volume. You can specify the following types: ■ IDE: `ide` ■ LSI Logic: `lsiLogic` ■ LSI Logic SAS: `lsiLogicsas` ■ BusLogic Parallel: `busLogic` ■ VMware Paravirtual SCSI: `paraVirtual`

Glance Images

9

In the OpenStack context, an image is a file that contains a virtual disk from which you can install an operating system on a virtual machine. 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 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. You can also use the `glance-import` tool to convert RAW, QCOW2, VDI, and VHD images to the VMDK format and use them in OpenStack.

This chapter includes the following topics:

Importing Images to the Image Service

You can use CLI commands or the VMware Integrated OpenStack dashboard to import images.

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 to VMDK before the import process (RAW, QCOW2, VDI, VHD).

Import Images Using the GUI

You can import images in the VMware Integrated OpenStack dashboard.

The following image formats are supported:

- VMDK
- ISO
- OVA

To upload images in another format, see Import Images in Unsupported Formats.

Note ISO images cannot be used to create volumes.

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** and click **Create Image**.
4. Configure the image.

Option	Action
Image Name	Enter a name for the image.
Image Description	Enter a description for the image.
Source Type	Select **File** to select a local file or **URL** to specify a remote file.
Format	Select **ISO** or **VMDK**. For images in OVA format, select **VMDK** as the disk format.
Disk Adapter Type	For VMDK images, select the adapter type.
Minimum Disk (GB)	Specify the minimum disk size for the image in gigabytes.
Minimum RAM (MB)	Specify the minimum RAM for the image in megabytes.
Visibility	Select **Public** to make the image available to all projects or **Private** to make the image available only to the current project.
Protected	Select **Yes** to prevent the image from being deleted.

5. (Optional) Click **Next** and configure metadata for the image.
6. Click **Create Image**.

What to do next

Tenants can launch OpenStack instances using the imported image. For instructions, see "Start an OpenStack Instance from an Image" in the *VMware Integrated OpenStack User's Guide*.

In the **Actions** column next to an image, you can edit the image, update its metadata, delete the image, or create a volume from the image.

Import Images in Supported Formats Using the CLI

You can import images using the command-line interface on the OpenStack Management Server.

The following image formats are supported:

- VMDK
- ISO
- OVA

To upload images in another format, see Import Images in Unsupported Formats.

Note ISO images cannot be used to create volumes.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Switch to the `root` user and load the cloud administrator credentials file.

```
sudo su -
source ~/cloudadmin.rc
```

3 Create the image in Glance.

```
glance image-create --name image-name --file image-path --disk-format {vmdk | iso} --
container_format bare --visibility {public | private} [--property vmware_adaptertype="adapter-
type"] [--property vmware_disktype="disk-type"] [--property vmware_ostype="operating-system"]
```

Option	Description
--name	Enter a name for the image file in the image service.
--file	Enter the path to the desired image file.
--disk_format	Enter the disk format of the source image. You can specify `iso` or `vmdk`. For images in OVA format, use `vmdk` as the disk format.
--container_format	Enter **`bare`**. The container format argument is not currently used by Glance.
--visibility	Enter **`public`** to make the image available to all users or **`private`** to make the image available only to the current user.
--property vmware_adaptertype	Specify the adapter type of the VMDK disk. If you do not include this parameter, the adapter type is determined by introspection. **Note** ■ For disks using paravirtual adapters, include this parameter and set it to **`paraVirtual`**. ■ For disks using LSI Logic SAS adapters, include this parameter and set it to **`lsiLogicsas`**.
--property vmware_disktype	Specify **`sparse`**, **`preallocated`**, or **`streamOptimized`**. If you do not include this parameter, the disk type is determined by introspection.
--property vmware_ostype	Specify the operating system on the image.

What to do next

You can run the `glance image-list` command to see the name and status of the images in your deployment.

Tenants can launch OpenStack instances using the imported image. For instructions, see "Start an OpenStack Instance from an Image" in the *VMware Integrated OpenStack User's Guide*.

Import Images in Unsupported Formats

You can use the `glance-import` tool to convert RAW, QCOW2, VDI, and VHD source images to the VMDK format.

You can also use this procedure to import images in the supported OVA and VMDK formats if desired.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Log in to the controller node as `viouser`.
3. Switch to the `root` user and load the cloud administrator credentials file.

   ```
   sudo su -
   source ~/cloudadmin.rc
   ```

4. Import the image.

   ```
   glance-import import --name image-name --url image-url --image-format format
   ```

Parameter	Description
--name	Enter a name for the image file in the image service.
--url	Enter the URL where the source image is located.
--image-format	Specify the format of the source image file. Non-VMDK images are converted automatically to the VMDK format. You can use the following formats: ■ VMDK ■ OVA ■ RAW ■ QCOW2 ■ VDI ■ VHD

   The task information and status is displayed. Large images might take some time to import. You can run the following command to check the status of the import task:

   ```
   glance task-show task-id
   ```

What to do next

You can run the `glance image-list` command to see the name and status of the images in your deployment.

Tenants can launch OpenStack instances using the imported image. For instructions, see "Start an OpenStack Instance from an Image" in the *VMware Integrated OpenStack User's Guide*.

Import a Virtual Machine Template as an Image

You can add virtual machine templates to your VMware Integrated OpenStack deployment as Glance images.

Prerequisites

- Verify that the virtual machine template is located in the same vCenter Server instance as your VMware Integrated OpenStack deployment.
- Verify that the virtual machine template does not have multiple disks, a CD-ROM drive, or a floppy drive.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. Switch to the `root` user and load the cloud administrator credentials file.

   ```
   sudo su -
   source ~/clouddadmin.rc
   ```

3. Create an empty Glance image in VMDK format.

   ```
   glance image-create --name image-name --disk-format vmdk --container-format bare
   ```

4. Add the location of the virtual machine template to the image.

   ```
   glance location-add image-uuid --url vi://vcserver-ip/datacenter-name/vm/folder-name/template-name
   ```

   You can check the **VM and Templates View** in the vSphere Client to confirm the location of the template.

The specified virtual machine template is imported as an image. You can launch OpenStack instances from the image or configure additional settings, such as image metadata.

Migrate an Image

You can migrate an image to another datastore while preserving its UUID and metadata.

Prerequisites

Determine the UUID of the image that you want to migrate and of the project containing the image. You can use the `openstack image list` command to display the UUID of each image and the `openstack image show` command to display the UUID of the project that contains a specified image.

Procedure

1. In the vSphere Web Client, open the **VMs and Templates** view and locate the image that you want to migrate.

   The image is located in the folder for the project that contains it.

2 Right-click the image and select **Clone to Template**.

3 Enter a new name for the image and click **Next**.

4 Select the desired compute resource and click **Next**.

5 Select the desired datastore and click **Next**.

6 Click **Finish**.

7 Record the name of the original image as shown in vSphere.

8 Delete the original image.

9 Rename the cloned image to the name of the original image.

The image is moved to the new datastore. You can continue to launch instances from it normally.

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 Server configured for VMware Integrated OpenStack. You can modify this behavior so that snapshots are stored as stream-optimized VMDK disks instead.

Before VMware Integrated OpenStack 2.5, the default behavior was to store Nova snapshots as stream-optimized 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 2.5, the default behavior was to store the Glance images 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 `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.

Supported Image Metadata

Image metadata is used for advanced configuration of Glance images. VMware Integrated OpenStack exposes additional capabilities through image metadata.

Note Configuring virtual interface quotas is not supported in NSX-T Data Center. The following metadata cannot be used with NSX-T Data Center deployments:

- `quota_vif_limit`
- `quota_vif_reservation`
- `quota_vif_shares_level`
- `quota_vif_shares_share`

If an image metadata and flavor extra spec conflict, the image metadata takes precedence over the flavor extra spec.

Table 9-1. Image Metadata in VMware Integrated OpenStack

Extra Spec	Description
`quota_cpu_limit`	Specify the maximum CPU allocation in MHz. The value `0` indicates that CPU usage is not limited.
`quota_cpu_reservation`	Specify the guaranteed CPU allocation in MHz.
`quota_cpu_reservation_percent`	Specify the guaranteed CPU allocation as a percentage of the actual CPU speed of the instance. This parameter takes precedence over the `cpu_reservation` parameter.
`quota_cpu_shares_level`	Specify the level of CPU shares allocated. You can enter **`custom`** and add the `cpu_shares_share` parameter to provide a custom value.
`quota_cpu_shares_share`	Specify the number of CPU shares allocated. If the `cpu_shares_level` parameter is not set to `custom`, this value is ignored.
`quota_disk_io_limit`	Specify the maximum disk transaction allocation in IOPS. The value `0` indicates that disk transactions are not limited.
`quota_disk_io_reservation`	Specify the guaranteed disk transaction allocation in IOPS.
`quota_disk_io_shares_level`	Specify the level of disk transaction shares allocated. You can enter **`custom`** and add the `disk_io_shares_share` parameter to provide a custom value.
`quota_disk_io_shares_share`	Specify the number of disk transaction shares allocated. If the `disk_io_shares_level` parameter is not set to `custom`, this value is ignored.
`quota_memory_limit`	Specify the maximum memory allocation in MB. The value `0` indicates that memory usage is not limited.
`quota_memory_reservation`	Specify the guaranteed memory allocation in MB.

Table 9-1. Image Metadata in VMware Integrated OpenStack (continued)

Extra Spec	Description
quota_memory_reservation_percent	Specify the guaranteed memory allocation as a percentage of the actual memory of the instance. The value 100 indicates that guest memory is also fully reserved. This parameter takes precedence over the memory_reservation parameter.
quota_memory_shares_level	Specify the level of memory shares allocated. You can enter **custom** and add the memory_shares_share parameter to provide a custom value.
quota_memory_shares_share	Specify the number of memory shares allocated. If the memory_shares_level parameter is not set to custom, this value is ignored.
quota_vif_limit	Specify the maximum virtual interface bandwidth allocation in Mbps. The value 0 indicates that virtual interface bandwidth is not limited.
quota_vif_reservation	Specify the guaranteed virtual interface bandwidth allocation in Mbps.
quota_vif_shares_level	Specify the level of virtual interface bandwidth shares allocated. You can enter **custom** and add the vif_shares_share parameter to provide a custom value.
quota_vif_shares_share	Specify the number of virtual interface bandwidth shares allocated. If the disk_io_shares_level parameter is not set to custom, this value is ignored.
vmware_extra_config	Specify custom configurations in JSON format. For example, '{“acpi.smbiosVersion2.7”:“FALSE”}' .
vmware_latency_sensitivity_level	Specify the latency sensitivity level for virtual machines. Setting this key will adjust certain settings on virtual machines.
vmware_resource_pool	Specify the resource pool on which to place new instances. If the name of the project containing the instance matches the name of a resource pool in your environment, the instance is placed in that resource pool by default. Setting this parameter overrides the default behavior and forces the instance to be placed in the specified resource pool.
vmware_storage_policy	Specify the storage policy used for new instances. If Storage Policy-Based Management (SPBM) is not enabled, this parameter is ignored.
vmware_tenant_vdc	Specify the UUID of the tenant virtual data center in which to place instances.
vmware_vm_group	Specify the DRS VM group in which virtual machines will be placed. If the specified VM group does not exist, instances will fail to power on.

10 Backup and Recovery

You can back up your VMware Integrated OpenStack installation to ensure that you can recover from errors that may occur.

This chapter includes the following topics:

Back Up Your Deployment

You can make backups of your management server data and OpenStack database.

For information about backing up Cinder, see Configure the Backup Service for Cinder.

Prerequisites

Prepare an NFS server to store backup information.

Procedure

1. Log in to the OpenStack Management Server.
2. Use the `viocli backup` command to back up desired information.
   - Run the following command to back up management server data:

     ```
     sudo viocli backup mgmt_server nfs-host-ip:/directory
     ```

     Backup files are stored in a folder named `vio_ms_`*`yyyymmddhhmmss`*.
   - Run the following command to back up the OpenStack database:

     ```
     sudo viocli backup openstack_db nfs-host-ip:/directory
     ```

     Backup files are stored in a folder named `vio_os_db_`*`yyyymmddhhmmss`*.

What to do next

If an error occurs on your deployment, you can recover individual nodes or the entire deployment. To recover individual nodes, see Recover OpenStack Nodes. To restore your deployment, see Restore Your Deployment from a Backup.

Configure the Backup Service for Cinder

You can configure Cinder to back up volumes to a network file system (NFS) server.

Prerequisites

- Create a shared NFS directory dedicated to storing Cinder backups.
- 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.

Procedure

1. Log in to the OpenStack Management Server as `viouser`.
2. If your deployment is not using a `custom.yml` file, copy the template `custom.yml` file to the `/opt/vmware/vio/custom` directory.

   ```
   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. Uncomment the `cinder_backup_driver` parameter and set its value to `cinder.backup.drivers.nfs`.
5. Uncomment the `cinder_backup_share` parameter and set its value to the location of the shared NFS directory.

   Use the format *nfs-host:path*. For example, `192.0.2.100:/cinder`.
6. Uncomment the `cinder_backup_mount_options` parameter and set it to your version of NFS.

   For example, enter **`vers=4`** to support NFS version 4.
7. Deploy the updated configuration.

   ```
   sudo viocli deployment configure
   ```

   Deploying the configuration briefly interrupts OpenStack services.

You can now use the `cinder backup–create` command to back up your Cinder volumes.

Restore Your Deployment from a Backup

You can restore your VMware Integrated OpenStack management server and OpenStack database from a backup.

If you want to recover individual nodes, see Recover OpenStack Nodes.

Prerequisites

Verify that you have a backup of the management server and database available. See Back Up Your Deployment.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Restore the OpenStack Management Server data.

```
sudo viocli restore mgmt_server backup-folder nfs-host-ip
```

Option	Description
backup-folder	Enter the name of the backup folder for OpenStack Management Server data. These folders are in the format `vio_ms_`*`yyyymmddhhmmss`*.
nfs-host-ip	Specify the IP address of the NFS host where the backup folder is located.

3 Restore the OpenStack database.

```
sudo viocli restore openstack_db backup-folder nfs-host-ip
```

Option	Description
backup-folder	Enter the name of the backup folder for the OpenStack database. These folders are in the format `vio_os_db_`*`yyyymmddhhmmss`*.
nfs-host-ip	Specify the IP address of the NFS host where the backup folder is located.

The OpenStack Management Server and OpenStack database are restored to the state of the backups.

Recover OpenStack Nodes

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 command-line interface.

When you recover a VMware Integrated OpenStack node, it returns to the state of a newly deployed node.

Prerequisites

- If you want to recover all database nodes, you must have a backup of the OpenStack database. See Back Up Your Deployment.
- Ensure that the datastore has sufficient free space to contain the original and recovered nodes at the same time. The recovery process will delete the original node, but space for both nodes is temporarily required. To avoid this issue, you can power off and delete the existing node before recovering it.

Procedure

1 Log in to the OpenStack Management Server as `viouser`.

2 Recover OpenStack nodes by node or role.

To display the nodes in your deployment, use the `viocli show` command. The values shown in the **VM Name** and **Role** columns can be used to recover nodes.

a To recover a non-database node, run the following command:

```
sudo viocli recover {–n node1... | –r role1... [–n node1...]}
```

Option	Description
–n	Enter the names of the nodes to recover.
–r	Enter the names of the roles to recover. All nodes assigned to the specified role will be recovered. You can specify –n in addition to this parameter to recover single nodes outside of the specified role.

b To recover a database node, run the following command:

```
sudo viocli recover {–n node1... | –r role} –dn backup–name –nfs nfs–host:/backup–folder
```

Option	Description
–n	Enter the names of the database nodes to recover. You can specify `DB` nodes for HA deployments or the `ControlPlane` node for compact or tiny deployments.
–r	Specify `DB` for HA deployments or `ControlPlane` for compact or tiny deployments. All database nodes will be recovered.
–dn	Enter the folder containing the OpenStack database backup. OpenStack database backup folders are in `vio_os_db_`*yyyymmddhhmmss* format.
–nfs	Specify the NFS host and directory where the backup is located in the format *remote–host*:/*remote–dir*.

The recovery process may take several minutes. You can check the status of your node by viewing your OpenStack deployment in the vSphere Web Client.

Troubleshooting VMware Integrated OpenStack

11

If errors occur, you can perform troubleshooting actions to restore your OpenStack deployment to operating status.

This chapter includes the following topics:

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.

OpenStack Management Server Logs

Name and Location	Description
`/var/log/apache2/access.log`	Logs access to the VMware Integrated OpenStack Manager.
`/var/log/apache2/error.log`	Logs access errors for the VMware Integrated OpenStack Manager.
`/var/log/column/ansible.log`	Logs Ansible service activity.
`/var/log/jarvis/jarvis.log`	Logs Jarvis service activity.
`/var/log/jarvis/pecan.log`	Logs Pecan framework service activity.
`/var/log/oms/oms.log`	Logs VMware Integrated OpenStack Manager service activity.
`/var/log/oms/register-plugin.log`	Logs VMware Integrated OpenStack plugin registration activity.
`/var/log/osvmw/osvmw-exceptions.log`	Logs exceptions to osvmw service.
`/var/log/osvmw/osvmw.log`	Logs osvmw service activity.
`/var/log/viocli/viocli.log`	Logs viocli (VMware Integrated OpenStack CLI) service activity.

Name and Location	Description
/var/log/viomon/viomon.log	Logs VMware Integrated OpenStack monitoring activity.
/var/log/viopatch/*.log	Logs upgrade and patching activity.
/var/log/bootsequence.log	Logs booting activity.

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 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 Load Balancer 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.

VMware Integrated OpenStack Performance Tuning

If the performance of your VMware Integrated OpenStack deployment deteriorates, you can adjust the settings for various VMware Integrated OpenStack components.

Because VMware Integrated OpenStack is deployed in many different environments, recommended values for performance parameters are not given. Adjust these parameters based on your environment and the resources available to you.

The parameters in the following table are located in the `custom.yml` file. You must run the `viocli deployment configure` command before your changes can take effect.

Table 11-1. VMware Integrated OpenStack Performance Tuning Parameters

Name	Default Value	Description	Usage
`nova_ram_allocation_ratio`	1.5	Allocation ratio of virtual memory to physical memory for CPU filters	Increase the value to address the following error in `nova-placement-api.log`: `InvalidAllocationCapacityExceeded: Unable to create allocation for 'MEMORY_MB' on resource provider`
`nova_cpu_allocation_ratio`	16	Allocation ratio of virtual CPUs to physical CPUs for CPU filters	Increase the value to address the following error in `nova-placement-api.log`: `InvalidAllocationCapacityExceeded: Unable to create allocation for 'VCPU' on resource provider`

Table 11-1. VMware Integrated OpenStack Performance Tuning Parameters (continued)

Name	Default Value	Description	Usage
`nova_disk_allocation_ratio`	0.0	Allocation ratio of virtual disk space to physical disk space for disk filters	Increase the value to address the following error in `nova-placement-api.log`: `InvalidAllocationCapacityExceeded: Unable to create allocation for 'DISK_GB' on resource provider`
`keystone_token_expiration_time`	7200	Time in seconds that a token remains valid	Increase the value to address the following error in various log files: `WARNING keystoneclient.middleware.auth_token [-] Authorization failed for token`
`haproxy_nova_compute_client_timeout`	1200s	Time in seconds that the load balancer waits for a response from Nova acting as a client	Increase these values to address the following error in `nova-compute.log`: `Exception during message handling: Gateway Time-out (HTTP 504)`
`haproxy_nova_compute_server_timeout`	1200s	Time in seconds that the load balancer waits for a response from Nova acting as a server	
`haproxy_cinder_client_lb_timeout`	300s	Time in seconds that the load balancer waits for a response from Cinder acting as a client	Increase these values to address the following error in `cinder-volume.log`: `VolumeBackendAPIException: Bad or unexpected response from the storage volume backend API`
`haproxy_cinder_server_lb_timeout`	300s	Time in seconds that the load balancer waits for a response from Cinder acting as a server	

Display the VMware Integrated OpenStack vApp

If the VMware Integrated OpenStack vApp does not appear in vSphere, you may need to perform various actions.

Problem

VMware Integrated OpenStack installed successfully, but the vApp is not displayed in vSphere.

Solution

1 In a browser, open `https://mgmt-server-ip:8443/VIO` and log in with the administrator credentials for your vCenter Server instance.

2 If the status indicator is red, perform the following steps:

   a Click **Fix**.

   b Verify the certificate and click **OK**.

   c Log out of the vSphere Web Client and log in again.

3 If the problem persists, confirm that the OpenStack Management Server can connect to the vCenter Server instance.

4 Log in to the OpenStack Management Server and check the logs in the `/var/log/oms` folder to confirm that the OpenStack Management Server service initiated properly.

5 Restart the OpenStack Management Server service.

```
service oms restart
```

6 Log out of the vSphere Web Client and log in again.

7 If the problem persists, log in to the vCenter Server virtual machine and restart the vSphere Web Client service.

   - For vSphere 6.5 or later, run the following commands:

```
service-control --stop vsphere-client
cd /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
rm -rf *
cd /usr/lib/vmware-vsphere-client/server/work
rm -rf *
service-control --start vsphere-client
```

   - For vSphere 6.0, run the following commands:

```
service vsphere-client stop
cd /etc/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
rm -rf *
service vsphere-client start
```

   - For vSphere 5.5, run the following commands:

```
service vsphere-client stop
cd /var/lib/vmware/vsphere-client/vc-packages/vsphere-client-serenity/
rm -rf * service vsphere-client start
```

8 Log out of the vSphere Web Client and log back in.

Resynchronize Availability Zones

In an environment with multiple vCenter Server instances, the names of availability zones on the VMware Integrated OpenStack dashboard might differ from those on the OpenStack Management Server.

If you use the command-line interface to rename availability zones, you might see different names in the vSphere Web Client and the VMware Integrated OpenStack dashboard. In the **Availability Zones** column on the **Manage > Nova Compute** tab for your deployment, desynchronized availability zones are displayed in red. You can resynchronize the availability zones to fix the issue.

Procedure

1 Log in to the OpenStack Management Server and list the availability zones in your OpenStack deployment.

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

2 Synchronize availability zones.

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

Troubleshoot Cinder Volume Backup Failure with Memory Error

Creating a backup of a Cinder volume on an NFS share fails with a memory error.

Problem

Attempting to create backup of a Cinder volume results in an out of memory error.

Cause

There is lack of available memory on the controller.

Solution

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 Depending on the available controller memory, reduce the value of the `cinder_backup_file_size` parameter.

b Change **=** symbol following the `cinder_backup_file_size` parameter to **:**.

For example,

```
cinder_backup_file_size: 52428800
```

3 Save the `custom.yml` file.

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

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

Troubleshoot Cinder Volume Backup Failure with Permission Denied Error

The first attempt to create a test backup of a Cinder volume on an NFS share fails with a permission denied error.

Problem

Attempting to verify the Cinder backup configuration results in a permission error when creating 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 controller 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 *
   ```

Solution

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

DCLI Cannot Connect to Server

If DCLI cannot connect to the OpenStack Management Server, you might need to restart the vAPI service.

Problem

When you start DCLI, the following error message is displayed:

```
ERROR: Unable to connect to the server.
```

Cause

DCLI cannot connect to the vAPI endpoint because the service is not running.

Solution

1. Log in to the OpenStack Management Server as `viouser`.

2 Check the status of the vAPI service.

```
sudo systemctl status vapi
```

The service is inactive.

```
vapi.service – VIO vAPI
   Loaded: loaded (/etc/systemd/system/vapi.service; disabled; vendor preset: enabled)
   Active: inactive (dead)
```

3 Restart the service.

```
sudo systemctl restart vapi
```

4 Check the status of the vAPI service again.

```
sudo systemctl status vapi
```

The service has restarted.

```
vapi.service – VIO vAPI
   Loaded: loaded (/etc/systemd/system/vapi.service; disabled; vendor preset: enabled)
   Active: active (running) since Wed 2018-06-27 04:46:00 UTC; 1s ago
  Process: 1983 ExecStartPre=/bin/mkdir -p /var/log/vmware/vapi (code=exited, status=0/SUCCESS)
 Main PID: 1985 (twistd)
   CGroup: /system.slice/vapi.service
           └─1985 /usr/bin/python /usr/bin/twistd --nodaemon --pidfile= -n web --port=9449 --wsgi
vmware.vapi.wsgi.application

Jun 27 04:46:00 vio-oms-01.mgt.sg.lab systemd[1]: Starting VIO vAPI...
Jun 27 04:46:00 vio-oms-01.mgt.sg.lab systemd[1]: Started VIO vAPI.
...
```

What to do next

Connect to the OpenStack Management Server again.

```
dcli +server http://mgmt-server-ip:9449/api +i
```

Using the OpenStack Management Server APIs

12

VMware Integrated OpenStack includes RESTful APIs that you can use to deploy and manage OpenStack.

Before using the APIs, you must authenticate with the OpenStack Management Server API endpoint using the administrator credentials for your vCenter Server instance. To authenticate, make a POST request to `https://mgmt-server-ip:8443/v1/j_spring_security_check` and include `j_username=vcenter-user&j_password=vcenter-password` in the request body.

After authentication, you are granted access to the APIs until the session expires. If using a web browser, you must accept the server certificate to establish a secure channel between the browser and the OpenStack Management Server before you can submit an API request.

For more information about APIs, see the VMware Integrated OpenStack API reference at https://code.vmware.com/apis/252. If you have installed VMware Integrated OpenStack, you can also view the API specifications at `https://mgmt-server-ip:8443/swagger-ui.html`.

VMware Integrated OpenStack Command Reference

13

VMware Integrated OpenStack includes the `viocli` utility to configure your deployment and the `viopatch` utility to manage and install patches. You run both command-line utilities on the OpenStack Management Server with `sudo`.

For NSX deployments, the `nsxadmin` utility is also provided to perform certain network-related operations. For more information, see the `nsxadmin` documentation at https://opendev.org/x/vmware-nsx/src/branch/master/doc/source/admin_util.rst.

The parameters supported by `viocli` and `viopatch` are described as follows. You can also run `viocli –h` or `viopatch –h` to display the supported parameters.

This chapter includes the following topics:

viocli backup Command

Use the `viocli backup` command to create a backup of either management server data or the OpenStack database. An NFS server must be available for VMware Integrated OpenStack to mount.

The `viocli backup` command uses the following syntax.

```
viocli backup {mgmt_server | openstack_db} [-d NAME] NFS-VOLUME
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
NFS-VOLUME	Mandatory	Name or IP address of the target NFS volume and directory in the format *remote-host*:*remote-dir*. For example: `192.168.1.77:/backups`

You can also run `viocli backup -h` or `viocli backup --help` to display the parameters for the command.

The backup file of the management server is labeled with a timestamp in `vio_ms_`*`yyyymmddhhmmss`* format. The backup file of the OpenStack database is labeled with a timestamp in `vio_os_db_`*`yyyymmddhhmmss`* format.

viocli certificate Command

Use the `viocli certificate` command to add, remove, and view certificates.

Note To generate a certificate signing request (CSR) or update an existing certificate, see viocli deployment Command.

The `viocli certificate` command supports a variety of actions to perform different tasks. The following parameters apply to all actions.

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-p` or `--progress`	Optional	Shows the progress of the current operation.

You can run `viocli certificate -h` or `viocli certificate --help` to display the actions and parameters for the command. You can also use the `-h` or `--help` option on any action to display parameters for the action. For example, `viocli certificate add -h` will show parameters for the `add` action.

The actions that `viocli certificate` supports are listed as follows.

viocli certificate add [-d *NAME*] --name *CERT-NAME* --cert *CERT-FILE* [-p]

Adds a certificate to the certificate store. The following additional parameters apply to the `add` action.

Parameter	Mandatory or Optional	Description
`--cert` *CERT-FILE*	Mandatory	Certificate to add. The certificate must be in PEM format.
`--name` *CERT-NAME*	Mandatory	Name of the certificate.

viocli certificate remove [-d *NAME*] --name *CERT-NAME* [-p]

Removes a certificate from the certificate store. The following additional parameters apply to the `remove` action.

Parameter	Mandatory or Optional	Description
`--name` *CERT-NAME*	Mandatory	Name of the certificate.

viocli certificate list [-d *NAME*] [--json JSON | -pretty PRETTY] [-p]

Lists all certificates in the certificate store. The following additional parameters apply to the `list` action.

Parameter	Mandatory or Optional	Description
`--json JSON` `--pretty PRETTY`	Optional	Displays output in JSON format or as formatted text. If you do not enter a value, `PRETTY` is used when the command is run interactively and `JSON` is used when the command is run noninteractively.

`viocli certificate show [–d NAME] ––name CERT–NAME [––json JSON | ––pretty PRETTY] [–p]`

Shows detailed information about a specified certificate. The following additional parameters apply to the `show` action.

Parameter	Mandatory or Optional	Description
`––name` *CERT-NAME*	Mandatory	Name of the certificate.
`––json JSON` `––pretty PRETTY`	Optional	Displays output in JSON format or as formatted text. If you do not enter a value, `PRETTY` is used when the command is run interactively and `JSON` is used when the command is run noninteractively.

viocli dbverify Command

Use the `viocli dbverify` command to check the VMware Integrated OpenStack database for problems such as duplicated or missing keys.

The `viocli dbverify` command uses the following syntax.

```
viocli dbverify [–d NAME]
```

Parameter	Mandatory or Optional	Description
`–d` *NAME* or `––deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.

You can also run `viocli dbverify –h` or `viocli dbverify ––help` to display the parameters for the command.

viocli deployment Command

Use the `viocli deployment` command to manage your VMware Integrated OpenStack deployment.

The `viocli deployment` command supports a variety of actions to perform different tasks. The following parameters apply to all actions.

Parameter	Mandatory or Optional	Description
`–d` *NAME* or `––deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`–p` or `––progress`	Optional	Shows the progress of the current operation.

You can run `viocli deployment –h` or `viocli deployment ––help` to display the parameters for the command. You can also use the `–h` or `––help` option on any action to display parameters for the action. For example, `viocli deployment configure –h` will show parameters for the `configure` action.

The actions that `viocli deployment` supports are listed as follows.

viocli deployment start [-d *NAME*] [-f] [-p]

Starts a deployment. The following additional parameters apply to the `start` action.

Parameter	Mandatory or Optional	Description
`-f` or `--force`	Optional	Force starts a deployment that is already running.

viocli deployment stop [-d *NAME*] [-p]

Stops a deployment.

viocli deployment pause [-d *NAME*] [-p]

Pauses a deployment.

viocli deployment resume [-d *NAME*] [-p]

Resumes a paused deployment.

viocli deployment reset_status [-d *NAME*] [-p]

Resets a deployment to running status.

Note Verify services before running this command.

viocli deployment configure [-d *NAME*] [--limit {controller | compute | db | memcache}] [--tags *TAGS*] [-p]

Updates the entire configuration for a deployment. The following additional parameters apply to the `configure` action.

Parameter	Mandatory or Optional	Description
`--limit {controller | compute | db | memcache}`	Optional	Updates the configuration for only the specified component.
`--tags` *TAGS*	Optional	Runs only those configuration tasks that are marked with the specified tags.

viocli deployment post-deploy [-d *NAME*] [-p]

Updates the post-deployment configuration.

viocli deployment run-custom-playbook [-d *NAME*] [-p]

Runs the custom Ansible playbook only.

viocli deployment cert-req-create [-d *NAME*] [-c *COUNTRY*] [-s *STATE*] [-l *CITY*] [-o *ORG*] [-u *ORG-UNIT*] [--hostname_list *HOST1*[,*HOST2*...]] [-p]

Creates a certificate signing request to send to a certificate authority. The following additional parameters apply to the `cert-req-create` action.

Parameter	Mandatory or Optional	Description
-c *COUNTRY* or --country_name *COUNTRY*	Optional	Two-letter ISO country code in which the organization applying for the certificate is located. If you do not include this option in the command, you will be prompted to enter a value.
-s *STATE* or --state_name *STATE*	Optional	Full name of the state or province. If you do not include this option in the command, you will be prompted to enter a value.
-l *CITY* or --locality_name *CITY*	Optional	Name of the town or city. If you do not include this option in the command, you will be prompted to enter a value.
-o *ORG* or --organization_name *ORG*	Optional	Legal name of the organization. If you do not include this option in the command, you will be prompted to enter a value.
-u *ORG-UNIT-NAME* or --organization_unit_name *ORG-UNIT-NAME*	Optional	Name of the department or organizational unit. If you do not include this option in the command, you will be prompted to enter a value.
--hostname_list *HOST1*[,*HOST2*...]	Optional	List of hostnames, separated with commas. If you do not include this option in the command, you will be prompted to enter a value.

viocli deployment cert-update [-d *NAME*] [-f *CERT-PATH*] [-p]

Updates the certificate used by VMware Integrated OpenStack. The following additional parameters apply to the `cert-update` action.

Parameter	Mandatory or Optional	Description
`-f` *CERT-PATH* or `--file` *CERT-PATH*	Optional	Absolute path to the desired certificate file. The certificate must be in PEM format.

viocli deployment getlogs [-d *NAME*] [--node *NODE*] [-svc *COMPONENT*] [-nrl] [--recent-logs] [-p]

Obtains log files for the current deployment, including executed Ansible commands and output. Log files are written to `/var/log/viocli/viocli.log` and rotated after they reach 100 MB. Only the most recent seven rotations are retained.

The following additional parameters apply to the `getlogs` action.

Parameter	Mandatory or Optional	Description
`--node` *NODE*	Optional	Obtains log files for the specified nodes only. The following values are supported: ■ `ceilometer` ■ `compute` ■ `controller` ■ `db` ■ `dhcp` ■ `lb` ■ `local` ■ `memcache` ■ `mongodb` ■ `mq` ■ `storage`
`-nrl` or `--non-rollover-log-only`	Optional	Collects only those logs that have not been archived.
`--recent-logs`	Optional	Collects only the log file to which the service process is currently writing.

viocli deployment default [-d *NAME*] [-p]

Returns the name of the default deployment.

viocli deployment status [-d *NAME*] [--period *SECONDS*] [--format {text | json}] [-p]

Assesses the status of a deployment in terms of the following:

- Synchronization problems between the management server and OpenStack nodes
- Connections to OpenStack processes and average connection count

- Interrupted network connections
- OpenStack database problems
- Missing processes

The following additional parameters apply to the `status` action.

Parameter	Mandatory or Optional	Description
`--period` *SECONDS*	Optional	Uses data from the specified period (in seconds) only. For example, `--period 300` will assess the status of the deployment in the last 5 minutes.
`--format {text | json}`	Optional	Outputs the status report in the specified format. If you do not enter a value, `text` is used by default.

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 [--verbose]
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
DC_NAME	Mandatory	Specifies a data center by name.
DS_NAME	Mandatory	Specifies a datastore by name.

You can also run `viocli ds-migrate-prep -h` or `viocli ds-migrate-prep --help` to display the parameters for the command.

viocli epops Command

Use the `viocli epops` command to manage the End Point Operations Management agent.

End Point Operations Management is a component of VMware vRealize Operations Manager. For more information, see the *vRealize Operations Manager Help* document for your version.

The `viocli epops` command supports a variety of actions to perform different tasks. The following parameters apply to all actions.

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.

You can run `viocli epops –h` or `viocli epops ––help` to display the parameters for the command. You can also use the `–h` or `––help` option on any action to display parameters for the action. For example, `viocli epops install –h` will show parameters for the `install` action.

The actions that `viocli epops` supports are listed as follows.

`viocli epops install [–d` *NAME*`] –s` *TGZ–FILE* `–c` *PROP–FILE*

Installs the End Point Operations Management agent. The following additional parameters apply to the `install` action.

Parameter	Mandatory or Optional	Description
`–s` *TGZ-FILE* or `––source` *TGZ-FILE*	Mandatory	Local path or URL to the agent installer package.
`–c` *PROP-FILE* or `––config` *PROP-FILE*	Mandatory	Local path to the agent configuration file.

`viocli epops uninstall [–d` *NAME*`]`

Uninstalls the End Point Operations Management agent.

`viocli epops reconfig [–d` *NAME*`] –c` *PROP–FILE*

Updates the configuration of the End Point Operations Management agent. The following additional parameters apply to the `reconfig` action.

Parameter	Mandatory or Optional	Description
`–c` *PROP-FILE* or `––config` *PROP-FILE*	Mandatory	Local path to the agent configuration file.

`viocli epops start [–d` *NAME*`]`

Starts the End Point Operations Management agent.

`viocli epops stop [–d` *NAME*`]`

Stops the End Point Operations Management agent.

viocli identity Command

Use the `viocli identity` command to configure Keystone for domains with AD or LDAP backends. The command calls the OpenStack Management Server API to store knowledge of Keystone domains and dictionary variables.

The `viocli identity` command supports a variety of actions to perform different tasks. The following parameters apply to all actions.

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-p` or `--progress`	Optional	Shows the progress of the current operation.

You can run `viocli identity -h` or `viocli identity --help` to display the parameters for the command. You can also use the `-h` or `--help` option on any action to display parameters for the action. For example, `viocli identity add -h` will show parameters for the `add` action.

The actions that `viocli identity` supports are listed as follows.

viocli identity add [-d *NAME*] [--type {AD | LDAP}] [-p]

Configures a new identity source. The following additional parameters apply to the `add` action.

Parameter	Mandatory or Optional	Description
`--type {AD | LDAP}`	Optional	Type of backend for the domain. If you do not include the `--type` parameter in the command, you will be prompted to enter the backend type.

viocli identity remove [-d *NAME*] --id *DOMAIN* [-p]

Removes an identity source from the list. The local (ID `0`) and default (ID `1`) domains cannot be removed.

The following additional parameters apply to the `remove` action.

Parameter	Mandatory or Optional	Description
`--id` *DOMAIN*	Mandatory	Identifier of an identity source. The local domain is represented by `0` and the default domain by `1`.

viocli identity edit [-d *NAME*] --id *DOMAIN* [-p]

Changes the settings of an existing identity source. The local domain (ID `0`) cannot be edited.

The following additional parameters apply to the `edit` action.

Parameter	Mandatory or Optional	Description
`--id` *DOMAIN*	Mandatory	Identifier of an identity source. The local domain is represented by `0` and the default domain by `1`.

viocli identity list [–d *NAME*] [––json JSON | ––pretty PRETTY] [–p]

Displays all configured domains with their ID numbers and backend types. The following additional parameters apply to the `list` action.

Parameter	Mandatory or Optional	Description
`––json JSON` `––pretty PRETTY`	Optional	Displays output in JSON format or as formatted text. If you do not enter a value, `PRETTY` is used when the command is run interactively and `JSON` is used when the command is run noninteractively.

viocli identity show [–d *NAME*] ––id *DOMAIN* [––json JSON | ––pretty PRETTY] [–p]

Displays detailed information about the specified domain. The following additional parameters apply to the `show` action.

Parameter	Mandatory or Optional	Description
`––id` *DOMAIN*	Mandatory	Identifier of an identity source. The local domain is represented by `0` and the default domain by `1`.
`––json JSON` `––pretty PRETTY`	Optional	Displays output in JSON format or as formatted text. If you do not enter a value, `PRETTY` is used when the command is run interactively and `JSON` is used when the command is run noninteractively.

viocli inventory-admin Command

Use the `viocli inventory–admin` command to compare the compute and block storage inventories against the vSphere inventory, discover and remove orphaned objects, and manage tenant virtual data centers.

Orphaned objects are defined as follows:

- Orphaned Nova instances are those for which a corresponding virtual machine does not exist in vSphere.
- Orphaned virtual machines are those for which a corresponding instance does not exist in the OpenStack database.
- Orphaned shadow virtual machines are those for which a corresponding Cinder volume does not exist in the OpenStack database.

The `viocli inventory-admin` command collects vCenter Server and OpenStack credentials from internal inventories. This command requires that you authenticate as an OpenStack administrator. The domain and user name of this account are set in `/root/cloudadmin.rc` as the `OS_PROJECT_DOMAIN_NAME`, `OS_USERNAME`, and `OS_USER_DOMAIN_NAME` variables. You can also set the password for this account as the `OS_PASSWORD` environment variable to avoid entering this password every time you run the command.

The `viocli inventory-admin` command supports a variety of actions to perform different tasks. The following parameters apply to all actions.

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`--json` `--pretty`	Optional	Displays output in JSON format or as formatted text. If you do not enter a value, `--pretty` is used when the command is run interactively and `--json` is used when the command is run noninteractively.
`--all`	Optional	Shows all objects instead of only orphaned objects.
`--force`	Optional	Runs the command without prompting for confirmation.
`--no-grace-period`	Optional	Ignores the grace period when determining whether objects are orphaned. Objects modified in the past 30 minutes are included in the results only when this parameter is set.

You can run `viocli inventory-admin -h` or `viocli inventory-admin --help` to display the parameters for the command. You can also use the `-h` or `--help` option on any action to display parameters for the action. For example, `viocli inventory-admin show-instances -h` will show parameters for the `show-instances` action.

The actions that `viocli inventory-admin` supports are listed as follows.

viocli inventory-admin show-instances [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Lists orphaned Nova instances.

viocli inventory-admin show-instance-vms [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Lists orphaned vSphere virtual machines.

viocli inventory-admin show-shadow-vms [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Lists orphaned shadow virtual machines.

viocli inventory-admin clean-instances [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Removes orphaned vSphere virtual machines.

viocli inventory-admin clean-instance-vms [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Removes orphaned vSphere virtual machines.

viocli inventory-admin clean-shadow-vms [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Removes orphaned shadow virtual machines.

viocli inventory-admin show-hypervisors [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Lists hypervisors with detailed information.

viocli inventory-admin show-availability-zones [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Lists availability zones and the hosts located in them.

viocli inventory-admin sync-availability-zones [-d *NAME*] [--filename *ZONE-MAP*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Synchronizes the availability zones in the environment with the specified map. The following additional parameters apply to the `sync-availability-zones` action.

Parameter	Mandatory or Optional	Description
`--filename` *ZONE-MAP*	Optional	Path to the file containing the availability zone map. The file must be in JSON format.

viocli inventory-admin create-tenant-vdc [-d *NAME*] --compute *COMPUTE-NODE* --name *VDC-NAME* --project-id *ID* [--cpu-reserve *CPU-MAX*] [--cpu-limit *CPU-MIN*] [--mem-reserve *MEMORY-MAX*] [--mem-limit *MEMORY-MIN*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Create a tenant virtual data center (VDC) with the specified settings. The following additional parameters apply to the `create-tenant-vdc` action.

Parameter	Mandatory or Optional	Description
`--compute` *COMPUTE-NODE*	Mandatory	Compute node on which to create the VDC.
`--name` *VDC-NAME*	Mandatory	Name of the tenant VDC.
`--project-id` *ID*	Mandatory	Project ID for the task.
`--cpu-reserve` *CPU-MIN*	Optional	CPU cycles in MHz to reserve for the VDC. If you do not enter a value, `0` is used by default.
`--cpu-limit` *CPU-MAX*	Optional	Maximum limit for CPU usage on the VDC (in MHz). If you do not enter a value, CPU usage is not limited.
`--mem-reserve` *MEMORY-MIN*	Optional	Memory in megabytes to reserve for the VDC. If you do not enter a value, `0` is used by default.
`--mem-limit` *MEMORY-MAX*	Optional	Maximum limit for memory consumption on the VDC (in megabytes). If you do not enter a value, memory consumption is not limited.

viocli inventory-admin list-tenant-vdcs [-d *NAME*] [--json | --pretty] [--all] [--force] [--no-grace-period]

Lists tenant VDCs.

viocli inventory-admin show-tenant-vdc [-d *NAME*] --id *ID* [--json | --pretty] [--all] [--force] [--no-grace-period]

Displays detailed information about the specified tenant VDC. The following additional parameters apply to the `show-tenant-vdc` action.

Parameter	Mandatory or Optional	Description
`--id` *ID*	Mandatory	Identifier of a tenant VDC.

`viocli inventory-admin delete-tenant-vdc [-d NAME] --id ID [--json | --pretty] [--all] [--force] [--no-grace-period]`

Deletes the specified tenant VDC. The following additional parameters apply to the `delete-tenant-vdc` action.

Parameter	Mandatory or Optional	Description
`--id` *ID*	Mandatory	Identifier of a tenant VDC.

viocli lbaasv2-enable Command

The `viocli lbaasv2-enable` command is no longer supported.

To enable LBaaS through the command line interface, see "Configure LBaaS Using the CLI" in the *VMware Integrated OpenStack User's Guide*.

viocli recover Command

Use the `viocli recover` command to recover nodes or groups of nodes.

Because most OpenStack nodes are stateless, you can recover them without a backup file. However, a backup file is necessary to recover OpenStack database nodes.

The `viocli recover` command uses the following syntax.

```
viocli recover [-d NAME] {-n NODE1... | -r ROLE1... [-n NODE1...]} [-dn BACKUP -nfs NFS-VOLUME]
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-n, --node` *NODE*	Mandatory unless `-r` is used.	Recovers one or more nodes. You can specify multiple nodes separated with commas. To display the nodes in your deployment, use the `viocli show` command. The values shown in the **VM Name** column can be used as arguments for this parameter. For example, the following command recovers two nodes from the specified NFS backup file. `viocli recover -n VIO-DB-0 VIO-DB-1 -dn vio_os_db_20150830215406 -nfs 10.146.29.123:/backups`
`-r` *ROLE* or `--role` *ROLE*	Mandatory unless `-n` is used.	Recovers all nodes assigned to the specified role. You can specify multiple roles separated with commas. You can also specify `-n` or `--node` in the same command to recover additional nodes that are not assigned to that role. To display the roles in your deployment, use the `viocli show` command. The values shown in the **Role** column can be used as arguments for this parameter. For example, the following command recovers the nodes assigned to the DB role from the specified NFS backup file. `viocli recover -r DB -dn vio_os_db_20150830215406 -nfs 10.146.29.123:/backups`

Parameter	Mandatory or Optional	Description
`-dn` *BACKUP* or `--dir-name` *BACKUP*	Mandatory for full OpenStack database recovery	Folder containing the OpenStack database backup files. OpenStack database backup folders are in `vio_os_db_`*`yyyymmddhhmmss`* format. This parameter is mandatory when recovering the following items: ■ For an HA deployment: the `DB` role or all three database nodes (`VIO-DB-0`, `VIO-DB-1`, and `VIO-DB-2`) ■ For a compact or tiny deployment: the `ControlPlane` role or the `VIO-ControlPlane-0` node
`-nfs` *NFS-VOLUME*	Mandatory for full OpenStack database recovery	Name or IP address of the target NFS volume and directory in the format *remote-host:/remote-dir*. For example: `192.168.1.77:/backups` This parameter is mandatory when recovering the following items: ■ For an HA deployment: the `DB` role or all three database nodes (`VIO-DB-0`, `VIO-DB-1`, and `VIO-DB-2`) ■ For a compact or tiny deployment: the `ControlPlane` role or the `VIO-ControlPlane-0` node

You can also run `viocli recover -h` or `viocli recover --help` to display the parameters for the command.

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 | openstack_db} [-d NAME] DIR-NAME NFS-VOLUME
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
DIR-NAME	Mandatory	Directory containing the backup file.
NFS-VOLUME	Mandatory	Name or IP address of the target NFS volume and directory in the format *remote-host:remote-dir*. For example: `192.168.1.77:/backups`

You can also run `viocli restore -h` or `viocli restore --help` to display the parameters for the command.

The backup file of the VMware Integrated OpenStack management server is labeled with a timestamp in `vio_ms_`*`yyyymmddhhmmss`* format. The backup file of the VMware Integrated OpenStack database is labeled with a timestamp in `vio_os_db_`*`yyyymmddhhmmss`* format.

viocli rollback Command

The `viocli rollback` command is no longer supported.

To roll back a recent patch, see "Roll Back a VMware Integrated OpenStack Patch" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

To revert from a recent upgrade, see "Revert to a Previous VMware Integrated OpenStack Deployment" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

viocli services Command

Use the `viocli services` command to start or stop all OpenStack services.

The `viocli services stop` command stops only the services running in your deployment. To stop the entire cluster, including virtual machines, run the `viocli deployment stop` command instead.

The `viocli services` command uses the following syntax.

```
viocli services [-d NAME] {start | stop}
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.

You can also run `viocli services -h` or `viocli services --help` to display the parameters for the command.

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 [-d NAME] [-i | -p] [--verbose]
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-i` or `--inventory`	Optional	Displays the contents of the inventory file for the current deployment.
`-p` or `--inventory-path`	Optional	Displays the path to the inventory file for the current deployment.
`--verbose`	Optional	Displays output in verbose mode.

To obtain a list of nodes, run `viocli show` without the `-i` or `-p` options.

You can also run `viocli show -h` or `viocli show --help` to display the parameters for the command.

viocli upgrade Command

Use the `viocli upgrade` command to upgrade between major versions of VMware Integrated OpenStack.

The `viocli upgrade` command supports a variety of actions to perform different tasks. The following parameters apply to all actions.

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-p` or `--progress`	Optional	Shows the progress of the current operation.

You can run `viocli upgrade -h` or `viocli upgrade --help` to display the parameters for the command. You can also use the `-h` or `--help` option on any action to display parameters for the action. For example, `viocli upgrade prepare -h` will show parameters for the `prepare` action.

`viocli upgrade mgmt_server [-d` *NAME*`]` *DIR-NAME NFS-VOLUME* `[-p]`

Upgrades the management server database and configuration to the desired version. The following additional parameters apply to the `mgmt_server` action.

Parameter	Mandatory or Optional	Description
DIR-NAME	Mandatory	Directory containing the backup file.
NFS-VOLUME	Mandatory	Name or IP address of the target NFS volume and directory in the format *remote-host*:*remote-dir*. For example: `192.168.1.77:/backups`

`viocli upgrade prepare [-d` *NAME*`]` *BLUE-OMS-SERVER NFS-DIR-NAME* `[`*BLUE-VIOUSER-PASSWORD*`] [-f] [-p]`

Prepares the NFS server for the OpenStack Management Server upgrade. The following additional parameters apply to the `prepare` action.

Parameter	Mandatory or Optional	Description
BLUE-OMS-SERVER	Mandatory	IP address of the old OpenStack Management Server.
NFS-DIR-NAME	Mandatory	Local mount point to attach the target NFS volume.
BLUE-VIOUSER-PASSWORD	Optional	Password of the `viouser` account on the old OpenStack Management Server. If you do not include this option in the command, you will be prompted to enter the password.
`-f` or `--force`	Optional	Runs the command without prompting for confirmation.

`viocli upgrade openstack [-d NAME] [-n NEW-DEPLOY] [-f] [-p]`

Upgrades the VMware Integrated OpenStack deployment to the desired version.

Note If possible, use the vSphere Web Client to upgrade your deployment instead of this command.

The following additional parameters apply to the `openstack` action.

Parameter	Mandatory or Optional	Description
`-n` *NEW-DEPLOY*	Optional	Name of the deployment for the new version. If you do not include this option in the command, you will be prompted to enter a name.
`-f` or `--force`	Optional	Runs the command without prompting for confirmation.

viocli volume-migrate Command

Use the `viocli volume-migrate` command to migrate one or more non-attached Cinder volumes from one datastore to another.

Note To migrate attached volumes, you must migrate the entire instance.

To migrate volumes for shadow virtual machines, use the `viocli ds-migrate-prep` command and then complete the migration using the vSphere Web Client.

The `viocli volume-migrate` command uses the following syntax.

```
viocli volume-migrate [-d NAME] {--volume-ids UUID1[,UUID2...] | --source-dc SRC-DC-NAME --source-ds
SRC-DS-NAME] DEST-DC-NAME DEST-DS-NAME [--ignore-storage-policy]
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`--volume-ids` *UUID1*	Mandatory unless `--source-dc` and `--source-ds` are used.	Migrates one or more volumes specified by UUID. To specify multiple volumes, separate the UUIDs with commas. For example, the following command migrates two volumes to datastore DS-01 in data center DC-01. `viocli volume-migrate --volume-ids 25e121d9-1153-4d15-92f8-c92c10b4987f,4f1120e1-9ed4-421a-b65b-908ab1c6bc50 DC-01 DS-01`
`--source-dc` *SRC-DC-NAME*	Mandatory unless `--volume-ids` is used.	Identifies the source data center. This option must be used together with the `--source-ds` option.
`--source-ds` *SRC-DS-NAME*	Mandatory unless `--volume-ids` is used.	Identifies the source datastore. This option must be used together with the `--source-dc` option. 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. `viocli volume-migrate --source-dc DC-01 --source-ds DS-01 DC-02 DS-02`
DEST-DC-NAME	Mandatory	Specifies the destination data center.

Parameter	Mandatory or Optional	Description
DEST-DS-NAME	Mandatory	Specifies the destination datastore.
`--ignore-storage-policy`	Optional	Ignores storage policy compliance check. This parameter enables volume migration when the destination datastore does not comply with the storage policy of the migrated volume.

You can also run `viocli volume-migrate -h` or `viocli volume-migrate --help` to display the parameters for the command.

viocli vros Command

Use the `viocli vros` command to enable VMware Integrated OpenStack to interoperate with vRealize Automation.

The `viocli vros` command uses the following syntax.

```
viocli vros enable [-d NAME] -vt VRA-TENANT -vh VRA-HOST -va VRA-ADMIN -vrh VROS-HOST
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-vt` *VRA-TENANT* or `--vra_tenant` *VRA-TENANT*	Mandatory	Tenant to which the vRealize Automation system administrator belongs.
`-vh` *VRA-HOST* or `--vra_host` *VRA-HOST*	Mandatory	IP or host name of vRealize Automation.
`-va` *VRA-ADMIN* or `--vra_admin` *VRA-ADMIN*	Mandatory	Username of the vRealize Automation system administrator.
`-vrh` *VROS-HOST* or `--vros_host` *VROS-HOST*	Mandatory	IP or host name for the vRealize Orchestrator OpenStack Plug-In service.

You can also run `viocli vros -h` or `viocli vros --help` to display the parameters for the command.

viopatch add Command

Use the `viopatch add` command to add new patches to your deployment so that you can install them.

The `viopatch add` command uses the following syntax.

```
viopatch add -l PATCH-LOCATION
```

Parameter	Mandatory or Optional	Description
`-l` *PATCH-LOCATION* or `--location` *PATCH-LOCATION*	Mandatory	Path of the patch file to add.

You can also run `viopatch add –h` or `viopatch add ––help` to display the parameters for the command.

viopatch install Command

Use the `viopatch install` command to install VMware Integrated OpenStack patches.

You must use the `viopatch add` command to add patches before you can install them.

The `viopatch install` command uses the following syntax.

```
viopatch install [–d NAME] –p PATCH–NAME –v PATCH–VERSION
```

Parameter	Mandatory or Optional	Description
`–d` *NAME* or `––deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`–p` *PATCH-NAME* or `––patch` *PATCH-NAME*	Mandatory	Name of the patch to install.
`–v` *PATCH-VERSION* or `––version` *PATCH-VERSION*	Mandatory	Version of the patch to install.

You can also run `viopatch install –h` or `viopatch install ––help` to display the parameters for the command.

viopatch list Command

Use the `viopatch list` command to display all VMware Integrated OpenStack patches that have been added.

You can also run `viopatch list –h` or `viopatch list ––help` to display the parameters for the command.

viopatch snapshot Command

Use the `viopatch snapshot` command to take and manage snapshots of your OpenStack deployment for pre-patch backup.

Important The `viopatch snapshot take` command stops OpenStack services. Services will be started again when the patch is installed. If you decide not to install a patch after taking a snapshot, you can manually start OpenStack services by running the `viocli services start` command.

The `viopatch snapshot` command uses the following syntax.

```
viopatch snapshot {take | revert | remove | list} [–d NAME] [–p]
```

Parameter	Mandatory or Optional	Description
`-d` *NAME* or `--deployment` *NAME*	Optional	Name of the deployment to use. If you do not enter a value, the default deployment is used.
`-p` or `--progress`	Optional	Shows the progress of the current operation.

You can also run `viopatch snapshot -h` or `viopatch snapshot --help` to display the parameters for the command.

viopatch uninstall Command

The `viopatch uninstall` command is no longer supported.

To roll back a recent patch, see "Roll Back a VMware Integrated OpenStack Patch" in the *VMware Integrated OpenStack Installation and Configuration Guide*.

viopatch version Command

Use the `viopatch version` command to display the current version of VMware Integrated OpenStack.

You can also run `viopatch version -h` or `viopatch version --help` to display the parameters for the command.