VMware Integrated OpenStack Administration Guide

Update 2
Modified on 13 NOV 2018
VMware Integrated OpenStack 5.0

vmware®

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

https://docs.vmware.com/

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

docfeedback@vmware.com

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

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.

VMware Technical Publications Glossary

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

Updated Information

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.
03 JUL 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:

Adding Capacity to an OpenStack Deployment

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

Add Compute Clusters to an OpenStack Deployment

You can add compute clusters to your VMware Integrated OpenStack deployment to increase CPU 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 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 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 and click the **Nova Storage** tab.

3 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 and click the **Glance Storage** tab.

3 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.

Use Profiling to Trace OpenStack Deployments

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

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

Procedure

1 Configure Tracing of OpenStack Services

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

2 Use Tracing of OpenStack Services

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

Configure Tracing of OpenStack Services

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

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

Prerequisites

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

Procedure

1 Modify the `custom.yml` file to enable tracing.

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

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

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

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

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

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

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

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

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

```
viocli deployment configure
```

Note Pushing the configuration briefly interrupts OpenStack services.

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

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

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

You can now use the profiling feature.

Use Tracing of OpenStack Services

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

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

Prerequisites

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

Procedure

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

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

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

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

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

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

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

Add IP Address Ranges to a Network

You can add IP address ranges to the networks in your deployment.

Important The management and API access networks cannot include more than 100 IP addresses each.

Procedure

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

   The **Networks** tab lists the management and API access network configurations, including their IP address ranges.
3. Right-click the network that you want to modify and select **Add IP Range**.
4. Specify the IP address range that you want to add and click **OK**.

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

Modify Network DNS Settings

You can modify the DNS settings for the management and API access networks.

Important Modifying the network DNS settings will briefly interrupt the network connection.

Procedure

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

   The **Networks** tab lists the management and API access network configurations, including their DNS servers.
3. Right-click the network that you want to modify and select **Change DNS**.
4. Specify the IP address of the primary and secondary DNS servers and click **OK**.

Change the Syslog Server

You can modify the syslog server IP address to send logs to a different server after deployment.

Procedure

1. In the vSphere Web Client, select **Home > VMware Integrated OpenStack**.
2. Click the **Manage** tab and click the **Settings** tab.
3. Click **Syslog Server** and click **Edit**.
4. Enter the IP address, port, and protocol of your vRealize Log Insight syslog server and click **OK**.

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 password, perform the following steps:

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

2 Change the password in NSX-T.

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 Click the **Manage** tab and click the **Settings** tab.

3 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.

OpenStack Instances in vSphere

The instances that you create in your VMware Integrated OpenStack deployment appear as virtual machines in your vCenter Server inventory. Many restrictions apply to how you manage and work with OpenStack VMs.

In most cases, you must manage OpenStack virtual machines on the VMware Integrated OpenStack dashboard or CLI rather than in the vSphere Web Client.

OpenStack Features Supported in vSphere

vSphere supports certain OpenStack features.

OpenStack Feature	Supported in vSphere
Launch	YES
Reboot	YES
Terminate	YES
Resize	YES
Rescue	YES
Pause	NO
Un-pause	NO
Suspend	YES
Resume	YES
Inject Networking Inject Networking is supported only when the following conditions are present: ■ With nova network in Flat mode ■ With Debian- or Ubuntu-based virtual machines ■ At boot time	YES
Inject File	NO
Serial Console Output	YES
RDP Console	NO
Attach Volume	YES
Detach Volume	YES
Live Migration	YES
Snapshot	YES
iSCSI	YES
Fibre Channel	YES Supported through vSphere datastores
Set Admin Pass	NO

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

VM Operations in OpenStack

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

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

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

vCenter Server Features Not Supported in the OpenStack API

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

- Adding a host to a cluster

  OpenStack cannot add a host to a cluster in vSphere.

- Putting a host into maintenance mode

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

- Resource Pools

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

- vSphere snapshots

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

Customize Horizon Logos

You can customize the logos that appear on the VMware Integrated OpenStack dashboard login page and in the upper left corner of other pages.

By default, the VMware corporate logo is used on the login page and in the upper left corner of each page on the dashboard. You can upload a custom graphic file to the OpenStack Management Server and configure it to display as your login or dashboard logo.

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. Use SCP to transfer your custom logo file or files to a temporary directory on the OpenStack Management Server virtual machine as the `viouser` account.

   ```
   scp your-logo-file viouser@mgmt-server-ip:/home/viouser/
   ```

2. Log in to the OpenStack Management Server virtual machine as the `viouser` account.

   ```
   ssh viouser@mgmt-server-ip
   ```

3. Create the `/opt/vmware/vio/custom/horizon` directory and move your logo file to that directory.

   ```
   sudo mkdir -p /opt/vmware/vio/custom/horizon
   sudo mv /home/viouser/your-logo-file /opt/vmware/vio/custom/horizon/
   ```

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

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

5. Open the `/opt/vmware/vio/custom/custom.yml` file in a text editor.
   - To configure a login logo, modify the following parameter:

     ```
     #horizon_logo_splash: "/opt/vmware/vio/custom/horizon/logo_file"
     ```

   - To configure a dashboard logo, modify the following parameter:

     ```
     #horizon_logo: "/opt/vmware/vio/custom/horizon/logo_file"
     ```

6 Delete the number sign (#) to uncomment the parameter that you want to enable. Then replace `logo_file` with the name of your custom graphic file.

```
horizon_logo_splash: "/opt/vmware/vio/custom/horizon/your-login-logo"
horizon_logo: "/opt/vmware/vio/custom/horizon/your-dash-logo"
```

You can enable one or both parameters.

7 Deploy the updated configuration.

```
sudo viocli deployment configure --tags horizon
```

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

Option	Description
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
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
```

Modify In-Flight Encryption Settings

You can change the cipher suites used by HAProxy and specify whether to encrypt in-flight data transferred between internal endpoints.

All public API endpoints in a VMware Integrated OpenStack deployment use TLS 1.2 encryption. For HA deployments, traffic between internal endpoints is also encrypted using TLS. Because the internal endpoints in a compact or tiny deployment are located on a single virtual machine, traffic between internal endpoints is not encrypted for those deployment types by default.

When internal in-flight encryption is enabled, HAProxy acts as a Layer 4 load balancer instead of a Layer 7 load balancer for internal API calls and Horizon traffic. To ensure strong encryption performance, the Apache HTTP server on each controller terminates TLS for each individual OpenStack service. The Apache server then forwards the request over a local loopback service to the back-end service, such as Nova, Neutron, or Cinder. HAProxy also re-encrypts the request when sending it to a back-end controller node over the internal 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. Modify encryption settings as desired.
   - To adjust the cipher suites, uncomment the `haproxy_ssl_default_bind_ciphers` parameter and set its value to the desired cipher suite.
   - To toggle TLS protection for internal endpoints, uncomment the `internal_api_protocol` parameter and set its value to **`https`** (TLS enabled) or **`http`** (TLS disabled).

5 Deploy the updated configuration.

```
sudo viocli deployment configure
```

Deploying the configuration briefly interrupts OpenStack services.

6 If you changed the value of the `internal_api_protocol` parameter, update the Keystone endpoint URL accordingly.

a In the vSphere Web Client, select **Administration > OpenStack**.

Note The HTML5 vSphere Client does not currently support this operation. Use the Flex-based vSphere Web Client.

b Select the **KEYSTONE** endpoint and click the **Edit** (pencil) icon.

c In the **Update Endpoint** section, change the URL to begin with `http` or `https` depending on your configuration.

d Enter the administrator password and click **Update**.

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 Using Flavor Metadata.

You can manage tenant virtual data centers using the `viocli` utility, vAPIs, or Data Center Command-Line Interface (DCLI). This procedure uses the `viocli` utility as an example. For information about vAPI or DCLI usage, see Using the Tenant Virtual Data Center vAPIs.

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 > Compute > 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 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 a Provider Network with NSX-V Networking

With NSX-V 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.
- 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.
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 NSX-T Networking

With NSX-T 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.

Option	Description
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 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-V Networking

With NSX-V 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 an External Network with NSX-T Networking

For NSX-T 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).

Option	Description
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 Neutron Availability Zone for an NSX-V Deployment

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 vSphere Distributed Switch. You can create a new distributed switch for the zone if desired.
- In NSX-V, 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.
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.

Create a Neutron Availability Zone for an NSX-T Deployment

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 a separate DHCP profile and metadata proxy server for each availability zone. Availability zones can share an edge cluster or use separate edge clusters.

For information about creating a DHCP profile, see Create a DHCP Server Profile in the *NSX-T Administration Guide*. For information about creating a metadata proxy server, see Add a Metadata Proxy Server in the *NSX-T 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 `nsxv3_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 `nsxv3_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.
metadata_proxy	Enter the name or UUID of the metadata proxy server for the availability zone.
dhcp_profile	Enter the name or UUID of the DHCP profile for the availability zone.
native_metadata_route	(Optional) Specify the route used for the metadata proxy service. Enter an IP address with prefix in CIDR notation.
dns_domain	(Optional) Enter the DNS domain for hostnames in the availability zone.
nameservers	(Optional) Enter one or more DNS servers to configure for DHCP binding entries.
default_overlay_tz	(Optional) Enter the name or UUID of the default overlay transport zone.
default_vlan_tz	(Optional) Enter the name or UUID of the default VLAN transport zone.
switching_profiles	(Optional) Enter the UUIDs of the switching profiles for the availability zone.
dhcp_relay_service	(Optional) Enter the name or UUID of the DHCP relay service for the availability zone.
default_tier0_router	(Optional) Enter the name or UUID of the default tier-0 router for the 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.

Create a Layer 2 Bridge with NSX-V

A Layer 2 bridge allows compute nodes on a VXLAN to communicate with a physical VLAN.

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 `cloudadmin.rc` 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-V 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.

Create a Layer 2 Bridge with NSX-T

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

Prerequisites

In NSX-T, 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 `cloudadmin.rc` file.

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

4 Create a logical Layer 2 gateway, specifying the UUID of the NSX-T 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.

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-V and NSX-T networks, only tenant networks can be transparent.

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

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

Manage NSX-V Edge HA

For NSX-V 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-V

For NSX-V 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.

Use N-VDS Enhanced Data Path Mode with OpenStack

For NSX-T deployments, you can create networks and ports backed by a transport zone supporting N-VDS enhanced data path mode.

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*.

An NSX-managed virtual distributed switch (N-VDS) can operate in enhanced data path mode to provide network performance improvements needed by NFV workflows. For more information, see "Enhanced Data Path" in the *NSX-T Installation Guide*.

Prerequisites

Create a separate availability zone for the N-VDS in enhanced data path mode. See Create a Neutron Availability Zone for an NSX-T Deployment.

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 `nsxv3_disable_port_security_for_ens` parameter and set its value to `true`.

   Port security is not supported with N-VDS enhanced data path mode.
5. Deploy the updated configuration.

   ```
   sudo viocli deployment configure ––limit controller
   ```

   Deploying the configuration briefly interrupts OpenStack services.

What to do next

When you create networks that consume N-VDS enhanced data path, specify the availability zone created for it.

Configure BGP Dynamic Routing for Your VMware Integrated OpenStack Deployment

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

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

Prerequisites

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

Procedure

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

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

2. Create a provider subnet pool.

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

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

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

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

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

4. Create the external VXLAN network.

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

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

5. Create the external VXLAN subnet.

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

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

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

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

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

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

7 Update the NSX Edges with BGP advertisement.

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

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

8 Update the NSX Edges with BGP neighbors.

Use the IDs of the edges that you created earlier.

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

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

9 Update your physical routers.

a Set AS value to **65000**.

b Set BGP neighbours to *name_GW-EDGE1* and *name_GW-EDGE2*.

c Set to advertise itself as dynamic gateway.

10 Create and configure the BGP Speaker.

a Create the BGP speaker.

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

b Create BGP peers.

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

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

c Add the BGP peer to the BGP speaker.

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

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

d Associate the speaker with the VXLAN network.

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

11 (Optional) Create BGP routers for tenants.

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

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

```
neutron net-create name_Tenant1_LS1

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

neutron net-create name_Tenant1_LS2

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

b Create a router with BGP configuration.

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

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

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

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

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

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

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 and NSX-V deployments.

- For NSX-T deployments, MAC learning in VMware Integrated OpenStack is provided by NSX-T MAC learning. For more information, see "Understanding MAC Management Switching Profile" in the *NSX-T Administration Guide* for your version of NSX-T.

- For NSX-V 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-V deployments, performance will be affected because vNICs that request promiscuous mode receive a copy of every packet.
- For NSX-V 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.
2. Disable port security and security groups on the port where you want to configure MAC learning.

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

3. Enable MAC learning on the port.

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

Add an NSX-T Backend to an NSX-V Deployment

If you have deployed VMware Integrated OpenStack with NSX-V networking, you can specify an NSX-T backend for certain projects in your deployment.

Important This process will update your `custom.yml` file or automatically generate a `custom.yml` file if the file does not exist in your environment. After running the `viocli enable-tvd` command, do not delete or overwrite `custom.yml` or your configuration will be discarded.

Prerequisites

- Deploy VMware Integrated OpenStack with NSX-V networking.
- Deploy NSX-T and obtain the following parameters:
  - IP address of the NSX Manager
  - Username and password to access the NSX Manager
  - Overlay transport zone
  - VLAN transport zone
  - Tier-0 router
  - DHCP profile
  - Metadata proxy server

Procedure

1 Create compute clusters for any projects that you want to use NSX-T and configure those clusters as transport nodes in your NSX-T environment.

A compute cluster cannot be part of an NSX-V and NSX-T deployment at the same time.

2 Log in to the VMware Integrated OpenStack manager and enable the TVD plugin.

```
sudo viocli enable-tvd --nsx-mgr manager-ip --nsx-user username --nsx-passwd password [--nsx-insecure {true | false}] [--nsx-ca-file ca-file] [--nsx-overlay-tz overlay-zone] [--nsx-vlan-tz vlan-zone] [--nsx-tier0-rt t0-router] [--nsx-dhcp-profile profile] [--nsx-md-proxy mdp-server]
```

Option	Description
--nsx-mgr	IP address of the NSX Manager of your NSX-T deployment.
--nsx-user	User name of the NSX Manager administrator.
--nsx-passwd	Password for the NSX Manager administrator.
--nsx-insecure {true | false}	Specifies whether to verify the certificate of the NSX Manager server. The default value is `true`.
--nsx-ca-file	CA bundle files to use in verifying the certificate of the NSX Manager server. This option is ignored if you include the `--nsx-insecure true` option.
--nsx-overlay-tz	Name or UUID of the default NSX-T overlay transport zone used for creating tunneled isolated Neutron networks.
--nsx-vlan-tz	Name or UUID of the default NSX-T VLAN transport zone used for bridging between Neutron networks if no physical network has been specified.
--nsx-tier0-rt	Name or UUID of the default tier-0 router used to connect to tier-1 logical routers and configure external networks.
--nsx-dhcp-profile	Name or UUID of the NSX-T DHCP profile used to enable native DHCP service.
--nsx-md-proxy	Name or UUID of the NSX-T metadata proxy server used to enable native metadata service.

3 Map existing projects to your NSX-T or NSX-V backend.

```
openstack project plugin create project-uuid --plugin {nsx-v | nsx-t}
```

Projects without a mapping use the NSX-V backend by default.

Authentication and Identity

4

In VMware Integrated OpenStack, authentication and identity management are provided by the Keystone component together with an identity provider such as VMware Identity Manager.

This chapter includes the following topics:

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. Do not use `default` or `local` as a Keystone domain.
Bind user	Enter the user name to bind to Active Directory for LDAP requests.

Option	Description
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 domain controllers, 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. Check the AD domain setting to filter out users of the same name as the service users in OpenStack such as nova or cinder. **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 https://docs.microsoft.com/en-us/windows/desktop/ADSI/search-filter-syntax.
Group tree DN	(Optional) Enter the search base for groups. The LDAP suffix is used by default.
Group filter	(Optional) Enter an LDAP search filter for groups.
LDAP admin user	If the Keystone identity provider is configured to work with OpenLDAP, enter the LDAP admin user.

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 the **Validate** button to confirm your settings.

Validation verifies that the admin user exists and that users are available in the user tree DN + filter search.

7 Click **OK**.

Configuring Federated Identity

Keystone is the identity service that provides identity, token, catalog, and policy services for use by services in the OpenStack family. Federated identity is the method used to establish trusts between identity providers and the services provided by an OpenStack Cloud.

VMware Integrated OpenStack supports three types of federated identity: Keystone federation, SAML2 federation, and vRealize Automation or VMware Identity Manager federation.

The following terms are commonly used when configuring federated identity.

Identity Provider (IdP)	Stores information about users and groups. The IdP provides authentication.
Service Provider (SP)	Provides a service to an end-user. The SP has protected resources.
Agent	Actor or browser that wants to access protected resources.
SAML assertion	Contains information about a user as provided by an IdP.

Configure Keystone to Keystone Federation

Keystone to Keystone Federation allows multiple OpenStack deployments to share the same identity source. It is useful for cross-region sites where one site is used as the identity source.

Keystone to Keystone configuration uses two VMware Integrated OpenStack deployments. A Keystone instance on the site B deployment authenticates directly with a Keystone instance on the site A deployment. The Keystone instance on site A authenticates with the SQL/LDAP.

Prerequisites

Verify that the public endpoint of the site B deployment can be reached from the internal network of the site A deployment.

Procedure

1 Log into the VMware Integrated OpenStack deployment on site A.

a Enable Keystone as an Identity Provider

```
viocli federation idp-metadata set
```

Enter input for prompts.

Prompt	Description	Sample Value
Lang	Language	en
Organization	Identity Provider organization	SAML identity Provider
Identity provider display name	Identity Provider display name	OpenStack SAML Identity Provider
Organization URL	Human readable name to identify the Identity Provider in Keystone and VMware Integrated OpenStack	https://www.vmware.com
Company name of contact person	Valid company name of the Identity Provider contact person	Example, Inc.
Given name of contact person	Valid given name of the Identity Provider contact person	John
Surname of contact person	Valid surname of the Identity Provider contact person	Doe
Email of contact person	Valid email address for Identity Provider contact person	john.doe@vmware.com
Telephone of contact person	Valid phone number for Identity Provider contact person	+1 800 555 0100
Type of contact	Type of contact person such as other, technical, support, administrative, billing	technical

b Trigger VMware Integrated OpenStack identity configuration.

```
viocli identity configure
```

Following configuration, expect a period of downtime to your VMware Integrated OpenStack deployment.

c Add a Keystone Service Provider for the Keystone instance running on site B.

```
viocli federation service-provider add
```

Enter input for prompts.

Prompt	Description	Sample Value
Service Provider name	Unique name to identify the Service Provider. Name must not include special characters or spaces that the URL cannot interpret.	keystone_180
Description	Human readable name to identify the Service Provider.	Keystone @ 192.168.112.180
Keystone address	Address of the keystone instance on site B	https://192.168.112.180:5000
Keystone IdP name	Value must match the name of the current Keystone Identity Provider specified in the Keystone Service Provider.	keystone_160

d Trigger VMware Integrated OpenStack identity configuration.

```
viocli identity configure
```

2 Log into the VMware Integrated OpenStack deployment on site B.

a Add a Keystone Identity Provider.

```
viocli federation identity-provider add
```

Enter input for prompts.

Prompt	Description	Sample Value
Identity provider type	Enter **keystone**. Value is case insensitive.	keystone
Identity provider name	Unique name to identify the Identity Provider. Name must not include special characters or spaces that the URL cannot interpret.	keystone_160
Identity provider display name	Human readable name to identify the Identity Provider in Horizon.	Keystone @ 192.168.112.160
Description	Human readable name to identify the Identity Provider in Keystone and VMware Integrated OpenStack.	Keystone @ 192.168.112.160
Keystone address to be federated.	Address of the Keystone instance that acts as an Identity Provider.	192.168.112.160
Enter the name of the domain that federated users associate with.	Name of the domain to which all federated users belong. If uncertain of the domain, enter **Default**. If it does not exist, VMware Integrated OpenStack creates the domain.	Default
Enter the name of the groups that federated users associate with (separated by commas ",").	Name of the groups to which all federated users belong. If using a customized mapping file, include all defined groups. If no mapping file exists, VMware Integrated OpenStack creates the groups within the domain.	Keystone Federated Users

b Trigger VMware Integrated OpenStack identity configuration.

```
viocli identity configure
```

Following configuration, expect a period of downtime to your VMware Integrated OpenStack deployment.

What to do next

If you do not want to use the default mapping, you can customize mapping. See Customize Mapping.

Configure Generic SAML2 Integration

Using generic SAML2, you can integrate OpenStack Management Server with any Identity Provider solution within your organization. Generic SAML2 authenticates directly with the identity service.

Since this integration does not automatically associate the Keystone Service Provider with your Identity Provider solution, contact your Identity Provider deployment administrator to collect information for mapping before configuring SAML2 integration.

SAML2 Federation uses a single OpenStack Management Server deployment.

Procedure

1 Add an Identity Provider on the OpenStack Management Server deployment.

```
viocli federation identity-provider add
```

Enter input for prompts.

Prompt	Description	Sample Value
Identity provider type	Enter **`saml2`**. Value is case insensitive.	saml2
Identity provider name	Unique name to identify the Identity Provider. Name must not include special characters or spaces that the URL cannot interpret.	vmware_saml
Identity provider display name	Human readable name to identify the Identity Provider in Horizon. Appears in the Horizon drop down menu.	VMware Generic SAML2
Description	Human readable name to identify the Identity Provider in Keystone and VMware Integrated OpenStack.	VMware Identity Manager @ vio-identity-manager.eng.vmware.com
address	Endpoint address of the vIDM deployment	vio-identity-manager.eng.vmware.com
vIDM admin user	User must have permission to list users.	admin
vIDM admin password		vmware
Enter the name of the domain that federated users associate with.	Name of the domain to which all federated users belong. If uncertain of the domain, enter **`Default`**. If it does not exist, VMware Integrated OpenStack creates the domain.	Default
Enter the name of the groups that federated users associate with (separated by commas ",").	Name of the groups to which all federated users belong. If using a customized mapping file, include all defined groups. If no mapping file exists, VMware Integrated OpenStack creates the groups within the domain.	ALL USERS, Federated Users

2 Configure the deployment

```
viocli identity configure
```

Following configuration, expect a period of downtime to your VMware Integrated OpenStack deployment.

What to do next

If you do not want to use the default mapping, you can customize mapping. See Customize Mapping.

Configure VMware Identity Manager Integration

If you have an existing vRealize Automation or VMware Identity Manager deployment, you can use VMware Identity Manager as an identity provider solution.

Note VMware Identity Manager users must authenticate using the VMware Integrated OpenStack dashboard. The OpenStack command-line interface is not supported for VMware Identity Manager.

Procedure

1 Add VMware Identity Manager as the Identity Provider on the OpenStack Management Server deployment.

```
viocli federation identity-provider add
```

Enter input for prompts.

Prompt	Description	Sample Value
Identity provider type	Enter **`vidm`**. Value is case insensitive.	vidm
Identity provider name	Unique name to identify the Identity Provider. Name must not include special characters or spaces that the URL cannot interpret.	vidm
Identity provider display name	Human readable name to identify the Identity Provider in Horizon. Appears in the Horizon drop down menu.	VMware Identity Manager
Description	Human readable name to identify the Identity Provider in Keystone and VMware Integrated OpenStack.	VMware Identity Manager @ vio-identity-manager.eng.vmware.com
vIDM endpoint address	Endpoint address of the vIDM deployment	https://vio-identity-manager.eng.vmware.com
vIDM admin user	User must have permission to list users.	admin
vIDM admin password		vmware
Do not verify certificates when establishing TLS/SSL connections.	Enter **`True`** or **`False`**. True disables certificate verification when establishing TLS/SSL connections.	True
vIDM Tenant	Tenant name to be used when registering the Keystone instance in vIDM. ■ If integrating with vIDM, the value can be left blank. ■ If integrating with vRA, enter **`vsphere.local`**.	(blank)
Enter the name of the domain that federated users associate with.	Name of the domain to which all federated users belong. If uncertain of the domain, enter **`Default`**. If it does not exist, VMware Integrated OpenStack creates the domain.	Default
Enter the name of the groups that federated users associate with (separated by commas ",").	Name of the groups to which all federated users belong. If using a customized mapping file, include all defined groups. If no mapping file exists, VMware Integrated OpenStack creates the groups within the domain.	(blank)

2 Configure the deployment

```
viocli identity configure
```

Following configuration, expect a period of downtime to your VMware Integrated OpenStack deployment.

What to do next

If you do not want to use the default mapping, you can customize mapping. See Customize Mapping.

Customize Mapping

For Keystone to Keystone federation and VMware Identity Manager integration, VMware Integrated OpenStack provides default mapping rules. If you want to customize mapping rules to fit your configuration, you can use the VMware Integrated OpenStack UI.

Prerequisites

Configure Keystone to Keystone federation or VMware Identity Manager integration.

Procedure

1 Log in to the VMware Integrated OpenStack dashboard as an administrator.

2 Under Federation, click **Mappings** to see the current mappings.

3 Click **Edit** to configure a mapping according to your needs.

For more information about mappings, see the Mapping Combinations for Federation in the OpenStack documentation.

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-V 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.

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-V Security Policies in OpenStack

You can enforce NSX-V 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-V 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-V. You can change the order in the vSphere Client on the **Security > Firewall** page under **Networking and Security**.

Prerequisites

Create the desired security policies in NSX-V.

Procedure

1. Log in to the OpenStack Management Server.
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-V 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.
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 If you want to use additional security groups with security policies, you can perform the following steps:

- To associate an NSX-V security policy with a new security group, create the group and run the following command:

```
neutron security-group-update --policy=policy-id security-group-id
```

- To migrate an existing security group to a security policy-based group, log in to the active controller and 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-id
```

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.

7 Log in to the active controller and grant NSX-V security policies higher priority than 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
```

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-V or VDS networking. Importing virtual machines is not supported for NSX-T deployments.
- Verify that the virtual machines that you want to import are in the same vCenter Server instance.

Procedure

1. In the vSphere Web Client, add the clusters containing the desired virtual machines as compute clusters in your VMware Integrated OpenStack deployment.
2. Log in to the OpenStack Management Server.
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 `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
      ```

4 Connect to the VMware Integrated OpenStack vAPI endpoint.

```
dcli +server https://<mgmt-server-ip>:9449/api +i
```

If you cannot connect to the server, see Troubleshoot Unable to 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.

Option	Description
`--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.
`--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"}]'`

Control the State of an Instance

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

Procedure

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

   Items that appear in red text are disabled.

Track Instance Use

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

Procedure

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

   The Overview page shows the usage summary and project-specific usage information. You can specify a period of time for the usage information. Optionally, you can download a CSV summary.
4. (Optional) Specify a period of time for reporting and click **Submit**.
5. (Optional) Click **Download CSV Summary** to download a report of the usage.

Enable Huge Page Support

To provide important or required performance improvements for some workloads, you can enable OpenStack huge page support for up to 1 GB per page. Huge pages are requested explicitly by using flavor extra specs or image metadata.

Prerequisites

- Verify that VMware Integrated OpenStack 5.0 or later is running.
- Verify that your deployment includes vSphere 6.7 or later.

Procedure

1 Add a flavor extra spec that requests huge pages with `hw` and `quota` properties.

```
$ openstack flavor set m1.large --property hw:mem_page_size=large
$ openstack flavor set m1.large --property quota:memory_reservation_percent=100
```

2 Create an OpenStack instance with the huge page flavor as in the following example.

```
$ openstack server create --flavor m1.large --image ubuntu foobar
```

3 Log in to the guest OS on the VMware Integrated OpenStack console.

Enabling huge pages is different for every OS. The following example shows how to enable persistent huge pages on a Linux host.

a To allocate huge pages at run time, modify `/etc/default/grub` to include some huge page parameters.

```
echo 'GRUB_CMDLINE_LINUX="default_hugepagesz=1G hugepagesz=1G hugepages=2
transparent_hugepage=never"' > /etc/default/grub
```

b Update the bootloader.

```
update-grub2
```

c Reboot the instance.

d Check that the instance is using huge pages.

```
grep "Huge" /proc/meminfo
```

The value for `Hugepagesize` should be 1 GB or less.

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.
2. Select the **admin** project from the drop-down menu in the title bar.
3. Select **Admin > Compute > Images**.
4. Create a new image or choose an existing image.
5. Click the down arrow next to the flavor that you want to use and select **Update Metadata**.
6. In the **Available Metadata** pane, expand **VMware Driver Options** and click the **Add** (plus sign) icon next to **DRS VM group**.
7. 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.

Using Affinity and Anti-Affinity to Place OpenStack Instances

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

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

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

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

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

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

Prerequisites

- Verify that the intended filter configuration does not conflict with any existing administrative configuration, such as DRS rules that manage instance placement on hosts.
- Verify that you are running VMware Integrated OpenStack version 2.0.x or later.
- Verify that VMware Integrated OpenStack is running.

- Verify that you are using a Python nova-client version 2.17.0.6 or later as required for the ServerGroup API. Go to http://docs.openstack.org/user-guide/common/cli_install_openstack_command_line_clients.html.

Procedure

1. Using SSH, log in to the nova-client.
2. (Optional) Obtain the ID of the image you will use to create the instance.

   You can use the `nova image-list` command to view the list of available images and their ID values.
3. (Optional) Obtain the ID of the flavor you will use to define the instance .

   You can use the `nova flavor-list` command to view the list of flavor definitions and their ID values.
4. Create a new server group with the intended policy.
   a. Create a server group with the affinity policy:

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

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

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

   In both case, the CLI returns the auto-generated server group UUID, name, and policy.
5. Launch a new instance, using the `--image`, `--flavor`, and `--hint` flags to apply the server group affinity policy .

   ```
   nova boot --image IMAGE_ID --flavor FLAVOR_ID --hint group=SERVER_GROUP_UUID INSTANCE_NAME
   ```
6. Confirm that the new rule and the server group instances appear and are running correctly in the VMware Integrated OpenStack deployment in vCenter Server.

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

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

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

Prerequisites

- Verify that the intended anti-affinity filter configuration does not conflict with any existing administrative configuration, such as DRS rules that manage instance placement on hosts.
- Verify that you are running VMware Integrated OpenStack version 2.0.x or later.

- Verify that VMware Integrated OpenStack is running.
- Verify that you are using a Python nova-client version 2.17.0.6 or later, as required for the ServerGroup API. Go to http://docs.openstack.org/user-guide/common/cli_install_openstack_command_line_clients.html.

Procedure

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

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

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

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

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

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

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

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

Configure QoS Resource Allocation for Instances Using Flavor Metadata

You can control the QoS resource allocations, such as limits, reservations, and shares, for CPU, RAM, disk IOPS, and virtual network interface (VIF) by modifying the metadata of the flavor used to create the instance. All instances subsequently created using the flavor inherit the metadata settings.

QoS resource allocation can also be specified by image metadata. In the event of a conflict, the image metadata configuration overrules the flavor metadata configuration. See Configure QoS Resource Allocation for Instances Using Image Metadata.

Prerequisites

- Requires VMware Integrated OpenStack version 2.0.x or greater.

- Requires vSphere version 6.0 or greater.
- Verify that VMware Integrated OpenStack is running in vSphere.
- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

1. Log in to the VMware Integrated OpenStack dashboard as a cloud administrator.
2. Select the admin project from the drop-down menu in the title bar.
3. Select **Admin > System > Flavors**.
4. (Optional) Create a flavor specific to the set of QoS resource allocations.

   You must create a custom flavor to contain the specific configuration. This leaves the original flavor configuration intact and available for other uses.
5. Select the flavor to modify.
6. In the Actions column of the image listing, click the down arrow and select **Update Metadata**.
7. In the column under Available Metadata, expand the **VMware Quota** tab.

   Note If the VMware Quota tab is not present, the related metadata properties might already be configured.
8. Click the plus sign (**+**) next to the VMware Quota metadata property you want to add.

Tip You can add all the options simultaneously by the clicking the plus sign (**+**) on the VMware Quota tab.

   In the column under Existing Metadata, the newly added metadata properties appear.
9. Configure the metadata properties.

Metadata Property	Description
Quota: CPU Limit	Applies the `quota:cpu_limit` metadata property. Specifies the upper limit for CPU allocation in MHz. This parameter ensures that the instance never uses more than the defined amount of CPU allocation. Enter **`0`** for unlimited CPU allocation.
Quota: CPU Reservation	Applies the `quota:cpu_reservation` metadata property. Specifies the guaranteed minimum CPU reservation in MHz. This parameter ensures that the instance has the reserved amount of CPU cycles available during resource contention.
Quota: CPU Shares Level	Applies the `quota:cpu_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota:cpu_shares_value` metadata property. See Quota: CPU Shares Value below.

Metadata Property	Description
Quota: CPU Shares Value	Applies the `quota:cpu_shares_value` metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the `quota:cpu_shares_level` metadata property to **`custom`**. Otherwise this property is ignored.
Quota: Disk IO Limit	Applies the `quota:disk_io_limit` metadata property. Specifies the upper limit for disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance never uses more than the defined amount of disk IOPS, and can be used to enforce a limit on the instance's disk performance. Enter **`0`** for unlimited IOPS.
Quota: Disk IO Reservation	Applies the `quota:disk_io_reservation` metadata property. Specifies the guaranteed minimum disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance receives the reserved amount of disk IOPS during resource contention.
Quota: Disk IO Shares Level	Applies the `quota:disk_io_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota:disk_io_shares_share` metadata property (Quota: Disk IO Shares Value).
Quota: Disk IO Shares Value	Applies the `quota:disk_io_shares_share` metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the `quota:disk_io_shares_level` metadata property to **`custom`**. Otherwise this property is ignored.
Quota: Memory Limit	Applies the `quota:memory_limit` metadata property. Specifies the upper limit for memory allocation in MB. This parameter ensures that the instance never uses more than the defined amount of memory. Enter **`0`** for unlimited memory allocation.
Quota: Memory Reservation	Applies the `quota:memory_reservation` metadata property. Specifies the guaranteed minimum memory reservation in MB. This parameter ensures that the instance receives the reserved amount of memory during resource contention.
Quota: Memory Shares Level	Applies the `quota:memory_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota:memory_shares_share` metadata property (Quota: Memory Shares Value).
Quota: Memory Shares Value	Applies the `quota:memory_shares_share` metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the `quota:memory_shares_level` metadata property to **`custom`**. Otherwise this property is ignored.
Quota: VIF Limit	Applies the `quota:vif_limit` metadata property. Specifies the upper limit for VIF bandwidth in Mbps. This parameter ensures that the VIF never uses more than the defined amount of bandwidth. Enter **`0`** for unlimited bandwidth allocation.

Metadata Property	Description
Quota: VIF Reservation	Applies the `quota:vif_reservation` metadata property. Specifies the guaranteed minimum bandwidth for VIF in Mbps. This parameter ensures that the virtual adapter on the instance gets the reserved amount of bandwidth during resource contention. If the instance uses less than the reserved amount, the remainder is available to other virtual adapters.
Quota: VIF Shares Level	Applies the `quota:vif_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota:vif_shares_share` metadata property (Quota: VIF Shares Value).
Quota: VIF Shares Value	Applies the `quota:vif_shares_share` metadata property. in the event that 'custom' is used, this is the number of shares.

10 Click **Save**.

The flavor metadata is now configured for limits, reservations, and shares for CPU, IOPS, memory, and network bandwidth. This configuration is applied to all future OpenStack instances that are created from this flavor.

Configure QoS Resource Allocation for Instances Using Image Metadata

You can control the QoS resource allocations, such as limits, reservations, and shares, for CPU, RAM, disk IOPS, and virtual network interface (VIF) by modifying the metadata of the source image used to create the instance. All instances subsequently created from the image inherit the metadata settings.

QoS resource allocation for an instance can also be specified by flavor metadata. In the event of a conflict, the image metadata configuration overrules the flavor metadata configuration. See Configure QoS Resource Allocation for Instances Using Flavor Metadata.

Prerequisites

- Requires VMware Integrated OpenStack version 2.0.x or later.
- Requires vSphere version 6.0 or later.
- Verify that VMware Integrated OpenStack is running in vSphere.
- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.

Procedure

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

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

3 Select **Admin > System > Images**.

4 Click the image to modify.

5 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.

6 In the column under Available Metadata, expand the **VMware Quota** tab.

Note If the **VMware Quota** tab is not present, the related metadata properties might already be configured.

7 Click the plus sign (**+**) next to the VMware Quota metadata property you want to add.

Tip You can add all the options simultaneously by the clicking the plus sign (**+**) on the **VMware Quota** tab.

In the column under Existing Metadata, the newly added metadata properties appear .

8 Configure the metadata properties.

Metadata Property	Description
Quota: CPU Limit	Applies the `quota_cpu_limit` metadata property. Specifies the upper limit for CPU allocation in MHz. This parameter ensures that the instance never uses more than the defined amount of CPU allocation. Enter **`0`** for unlimited CPU allocation.
Quota: CPU Reservation	Applies the `quota_cpu_reservation` metadata property. Specifies the guaranteed minimum CPU reservation in MHz. This parameter ensures that the instance has the reserved amount of CPU cycles available during resource contention.
Quota: CPU Shares Level	Applies the `quota_cpu_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the `quota_cpu_shares_value` metadata property. See Quota: CPU Shares Value below.
Quota: CPU Shares Value	Applies the `quota_cpu_shares_value` metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the `quota_cpu_shares_level` metadata property to **`custom`**. Otherwise this property is ignored.
Quota: Disk IO Limit	Applies the `quota_disk_io_limit` metadata property. Specifies the upper limit for disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance never uses more than the defined amount of disk IOPS, and can be used to enforce a limit on the instance's disk performance. Enter **`0`** for unlimited IOPS.
Quota: Disk IO Reservation	Applies the `quota_disk_io_reservation` metadata property. Specifies the guaranteed minimum disk transactions in I/O operations per second (IOPS) in seconds. This parameter ensures that the instance receives the reserved amount of disk IOPS during resource contention.
Quota: Disk IO Shares Level	Applies the `quota_disk_io_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota_disk_io_shares_share` metadata property (Quota: Disk IO Shares Value).

Metadata Property	Description
Quota: Disk IO Shares Value	Applies the `quota_disk_io_shares_share` metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the `quota_disk_io_shares_level` metadata property to **`custom`**. Otherwise this property is ignored.
Quota: Memory Limit	Applies the `quota_memory_limit` metadata property. Specifies the upper limit for memory allocation in MB. This parameter ensures that the instance never uses more than the defined amount of memory. Enter **`0`** for unlimited memory allocation.
Quota: Memory Reservation	Applies the `quota_memory_reservation` metadata property. Specifies the guaranteed minimum memory reservation in MB. This parameter ensures that the instance receives the reserved amount of memory during resource contention.
Quota: Memory Shares Level	Applies the `quota_memory_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota_memory_shares_share` metadata property (Quota: Memory Shares Value).
Quota: Memory Shares Value	Applies the `quota_memory_shares_share` metadata property. Specifies the number of shares allocated to the instance. Apply this property only if you set the `quota_memory_shares_level` metadata property to **`custom`**. Otherwise this property is ignored.
Quota: VIF Limit	Applies the `quota_vif_limit` metadata property. Specifies the upper limit for VIF bandwidth in Mbps. This parameter ensures that the VIF never uses more than the defined amount of bandwidth. Enter **`0`** for unlimited bandwidth allocation.
Quota: VIF Reservation	Applies the `quota_vif_reservation` metadata property. Specifies the guaranteed minimum bandwidth for VIF in Mbps. This parameter ensures that the virtual adapter on the instance gets the reserved amount of bandwidth during resource contention. If the instance uses less than the reserved amount, the remainder is available to other virtual adapters.
Quota: VIF Shares Level	Applies the `quota_vif_shares_level` metadata property. Specifies shares level which maps to the predefined numeric value of shares. If the **`custom`** level is selected, you must include the`quota_vif_shares_share` metadata property (Quota: VIF Shares Value).
Quota: VIF Shares Value	Applies the `quota_vif_shares_share` metadata property. in the event that ‘custom’ is used, this is the number of shares.

9 Click **Save**.

The image metadata is now configured for limits, reservations, and shares for CPU, IOPS, memory, and network bandwidth. This configuration is applied to all future OpenStack instances that are created from this image.

Apply QoS Resource Allocation to Existing Instances

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

Prerequisites

- Requires an OpenStack flavor with the desired QoS resource allocation settings. See Configure QoS Resource Allocation for Instances Using Flavor Metadata.
- Requires VMware Integrated OpenStack version 2.0.x or greater.
- Verify that VMware Integrated OpenStack is running in vSphere.

Procedure

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

   The resizing process may take a few minutes.

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

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 Deploy the updated configuration.

```
sudo viocli deployment configure
```

Deploying the configuration briefly interrupts OpenStack services.

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

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

10 Select **Admin > Compute > Flavors**.

11 Create a new flavor or choose an existing flavor.

12 Click **Update Metadata** to the right of the flavor.

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

14 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 > Compute > 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, expand **CPU Pinning** and click the **Add** (plus sign) icon next to **CPU Pinning policy**.

7 Set the value of `hw:cpu_policy` to **`dedicated`** 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 `cloudadmin.rc` 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-V networking. Deployments with NSX-T 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 the *vSphere Networking* document for your version.
  - To enable DirectPath I/O, see "Enable Passthrough for a Network Device on a Host" in the *vSphere Networking* document for your version.
- 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 `cloudadmin.rc` file.

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

3. Log in to the OpenStack Management Server.
4. Create a passthrough-enabled port.

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

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

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-V networking. Deployments with NSX-T 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 the *vSphere Networking* document for your version of vSphere.
  - To enable DirectPath I/O, see "Enable Passthrough for a Network Device on a Host" in the *vSphere Networking* document for your version of vSphere.
- 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 > Compute > 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 > Compute > 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.

Request GPU Shared Device for an OpenStack Instance

You can request a shared GPU device for an OpenStack instance by adding a GPU profile to your VMware Integrated OpenStack deployment and configuring a flavor extra spec to request the virtual GPU.

Prerequisites

Verify that the appropriate driver for your GPU device is installed on the ESXi host.

Procedure

1 Using SSH, log in to the VMware Integrated OpenStack Management Server.

2 Create the `custom.yml` file if it does not exist.

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

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

4 Specify the GPU profile and frame buffer size by editing the `custom.yml` file.

   a Edit the `nova_gpu_profile` value to specify the GPU profile for all compute nodes, for example:

   ```
   nova_gpu_profile: grid_p100–4a
   ```

   b Edit the `nova_profile_fb_size_kb` value to specify the GPU frame buffer size, for example:

   ```
   nova_profile_fb_size_kb: 4096
   ```

   c Save the `custom.yml` file.

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

   Refresh of the configuration briefly interrupts the OpenStack services.

   ```
   viocli deployment configure ––tags nova_api_config
   ```

6 Create a flavor extra spec that requests one virtual GPU.

   ```
   openstack flavor set vgpu_1 ––property "vmware:vgpu=1"
   ```

   VMware Integrated OpenStack supports one GPU per VM.

7 Create an OpenStack instance with the virtual GPU device.

   ```
   openstack server create ––flavor vgpu_1 ––image cirros–0.3.5–x86_64–uec ––wait test–vgpu
   ```

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.

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

Extra Spec	Description	Configurable by Image Metadata
`vmware:hw_version`	Specify the hardware version used to create images. In an environment with different host versions, you can use this key to place instances on the correct hosts.	No
`vmware:latency_sensitivity_level`	Specify the latency sensitivity level for virtual machines. Setting this key will adjust certain settings on virtual machines.	Yes
`vmware:storage_policy`	Specify the storage policy used for new instances. If Storage Policy-Based Management (SPBM) is not enabled, this parameter is ignored.	Yes
`vmware:tenant_vdc`	Specify the UUID of the tenant virtual data center in which to place instances.	Yes
`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.	Yes
`hw:vifs_multi_thread`	Specify **`true`** to provide each virtual interface with its own transmit thread.	No
`quota:cpu_limit`	Specify the maximum CPU allocation in MHz. The value `0` indicates that CPU usage is not limited.	Yes
`quota:cpu_reservation`	Specifies the guaranteed CPU allocation in MHz.	Yes

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

Extra Spec	Description	Configurable by Image Metadata
`quota:cpu_reservation_percent`	Specifies the guaranteed CPU allocation as a percentage of the actual CPU speed of the instance. This parameter takes precedence over the `cpu_reservation` parameter.	Yes
`quota:cpu_shares_level`	Specifies the level of CPU shares allocated. You can enter **`custom`** and add the `cpu_shares_share` parameter to provide a custom value.	Yes
`quota:cpu_shares_share`	Specifies the number of CPU shares allocated. If the `cpu_shares_level` parameter is not set to `custom`, this value is ignored.	Yes
`quota:memory_limit`	Specify the maximum memory allocation in MB. The value `0` indicates that memory usage is not limited.	Yes
`quota:memory_reservation`	Specify the guaranteed memory allocation in MB.	Yes
`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.	Yes
`quota:memory_shares_level`	Specifies the level of memory shares allocated. You can enter **`custom`** and add the `memory_shares_share` parameter to provide a custom value.	Yes
`quota:memory_shares_share`	Specifies the number of memory shares allocated. If the `memory_shares_level` parameter is not set to `custom`, this value is ignored.	Yes
`quota:disk_io_limit`	Specify the maximum disk transaction allocation in IOPS. The value `0` indicates that disk transactions are not limited.	Yes
`quota:disk_io_reservation`	Specify the guaranteed disk transaction allocation in IOPS.	Yes

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

Extra Spec	Description	Configurable by Image Metadata
`quota:disk_io_shares_level`	Specifies the level of disk transaction shares allocated. You can enter **`custom`** and add the `disk_io_shares_share` parameter to provide a custom value.	Yes
`quota:disk_io_shares_share`	Specifies the number of disk transaction shares allocated. If the `disk_io_shares_level` parameter is not set to `custom`, this value is ignored.	Yes
`quota:vif_limit`	Specify the maximum virtual interface bandwidth allocation in Mbps. The value `0` indicates that virtual interface bandwidth is not limited.	Yes
`quota:vif_reservation`	Specify the guaranteed virtual interface bandwidth allocation in Mbps.	Yes
`quota:vif_shares_level`	Specifies the level of virtual interface bandwidth shares allocated. You can enter **`custom`** and add the `vif_shares_share` parameter to provide a custom value.	Yes
`quota:vif_shares_share`	Specifies the number of virtual interface bandwidth shares allocated. If the `disk_io_shares_level` parameter is not set to `custom`, this value is ignored.	Yes

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 settings for volume encryption, corresponding vSphere storage profile, and default adapter type.

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 > Volume > 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 configure encryption for the volume type, perform the following steps:

Note You cannot add or update encryption when the volume type is used by a volume.

a In the **Actions** column, select **Create Encryption**.

b Configure encryption as follows.

Option	Description
Provider	Enter the class providing encryption.
Control Location	Select **front-end** or **back-end**.
Cipher	(Optional) Enter the encryption algorithm. If you do not enter a value, the default for the specified provider is used.
Key Size (bits)	(Optional) Enter the size of the encryption key in bits. If you do not enter a value, the default for the specified provider is used.

c Click **Create Volume Type Encryption**.

8 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**.

9 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**.

10 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 of 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 Cinder Volume Adapter Type

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

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

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

Procedure

1 Implement the `custom.yml` file.

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

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

   a Uncomment the `cinder_volume_default_adapter_type` parameter.

   b Change the setting with a custom value, for example **`lsiLogicsas`**.

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

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

3 Save the `custom.yml` file.

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

```
viocli deployment configure
```

Note Pushing the configuration briefly interrupts OpenStack services.

Configure the Volume Snapshot Format

The vSphere template is the default volume snapshot format for a vCenter Server. This allows VMware Integrated OpenStack to snapshot volumes that are attached to instances.

To change the snapshot format to the copy-on-write disk format used in VMware Integrated OpenStack 4.1 or earlier, you set the `cinder_vmware_snapshot_format` parameter using a `custom.yml` file.

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_vmware_snapshot_format` parameter.

b Set the value to **`COW`**.

```
cinder_vmware_snapshot_format: COW
```

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.

Migrating Volumes Between Datastores

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

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

Migrate All Volumes from a Specified Datastore

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

Prerequisites

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

Procedure

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

   ```
   sudo su –
   ```

3. Prepare the volume for migration.

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

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

Option	Description
`-d DEPLOYMENT`	Indicates the name of the VMware Integrated OpenStack deployment.
`DC_NAME`	Indicates the data center name.
`DS_NAME`	Indicates the datastore name.

4. Place the datastore in maintenance mode.

   See the vSphere product documentation.

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

Migrate Unattached Cinder Volumes

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

Prerequisites

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

Procedure

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

2 Switch to root user.

```
sudo su -
```

3 Migrate the volume.

```
viocli volume-migrate [-d NAME] {--volume-ids UUID1[,UUID2...] | --source-dc SRC-DC-NAME --source-ds SRC-DS-NAME] DEST-DC-NAME DEST-DS-NAME [--ignore-storage-policy]
```

For more information about this command, see the viocli volume-migrate Command documentation.

Migrate Attached Cinder Volumes

To migrate an attached Cinder volume to a different datastore, you must migrate the virtual machine that corresponds to the instance to which it is attached.

Prerequisites

Detach any snapshots that are attached to the volume.

Procedure

1 Log in to the OpenStack Management Server and prepare 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.

2 In the vSphere Web Client, locate the virtual machine that corresponds to the compute instance to which the volume is attached.

3 Use Storage vMotion to migrate the virtual machine to a different datastore.

The volume migrates to the new datastore, but only the disk of the shadow VM moves to the new datastore. The shadow VM remains on the old datastore with no disk.

4 (Optional) To fix the disk of the shadow VM, run a volume detach procedure.

The detach operation disconnects the volume from the instance. Failures to read or write from the volume might occur.

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. You can also import RAW, QCOW2, VDI, and VHD images, which are automatically converted to the VMDK format during the image creation process.

This chapter includes the following topics:

Import Images Using the GUI

You can import images in the VMware Integrated OpenStack dashboard.

The following image formats are supported:

- VMDK
- ISO
- OVA
- RAW
- QCOW2
- VDI

- VHD

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 > Compute > 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.
Image Source	Select the image file.
Format	Select **ISO** or **VMDK**. For images in OVA, RAW, QCOW2, VDI, or VHD formats, 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 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

- RAW
- QCOW2
- VDI
- VHD

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 `cloudadmin.rc` file.

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

3. Create the image in Glance.

```
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"
```

Option	Description
image-name	Enter the name of the source image.
--disk-format	Enter the disk format of the source image. You can specify `iso` or `vmdk`. For images in other formats, including OVA, RAW, QCOW2, VDI, or VHD, use `vmdk` as the disk format.
--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. 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 `openstack 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*.

Add a VM Template as an Image

You can add existing VM templates to your VMware Integrated OpenStack deployment as Glance images. This enables users to boot instances, create bootable block storage volumes, and other functions available to Glance images.

Prerequisites

- Verify that the existing VMs template resides in the same vCenter Server as your VMware Integrated OpenStack deployment.
- Verify that the following conditions do apply.
  - The VM template does not have multiple disks.
  - The VM template does not have a CD-ROM drive.
  - The VM template does not have a floppy disk drive.

Procedure

1. Prepare the VM template.

   Configure the metadata settings as necessary.

   - The `vmware_ostype` is required for Windows images, but optional for Linux images.
   - The `hw_vif_model` is recommended for specifying NIC type. Before defining this setting, confirm the correct NIC type for this image template. For example, if this setting is undefined, the instance is provisioned with the E1000 NIC by default. To ensure another NIC is provisioned, define this setting appropriately.

     For example, to provision the VMXNET3 NIC, the metadata definition is **`hw_vif_model=VirtualVmxnet3`**.
   - The following metadata settings are not required.
     - `vmware_adaptertype`
     - `vmware_disktype`
2. Log in to the OpenStack management cluster.

3 Run the `glance` command to obtain, define, and import the image.

```
glance image-create --name <NAME> \
       --disk-format vmdk --container-format bare
       --property vmware_ostype=ubuntu64Guest
       --property hw_vif_model=VirtualVmxnet3

glance location-add <glance_image_UUID> --url "vi://<vcenter-host>/<datacenter-path>/vm/<sub-
folders>/<template_name> IMAGE_ID"
```

The `location-add` command points to the inventory path for the VM template and can refer to either VM or host. For example:

```
"vi://<datacenter-path>/vm/<template_name>"
or
"vi://<datacenter-path>/host/<host_name>/<template_name>"
```

The `vm` and `host` keywords in the inventory path represent the **VM and Templates View** and **Host and Cluster View** hierarchy in vSphere.

Migrate an Existing Image

You can migrate images between datastores in a way that preserves their UUID and metadata.

Prerequisites

Verify that both the current and destination datastore are available.

Procedure

1 Using SSH, log in to the controller01 node.

2 Switch to root user.

```
sudo su -
```

3 View a list of images.

```
openstack image list
```

The result lists image UUIDs, names, and statuses.

```
+--------------------------------------+--------------------------------------+--------+
| ID                                   | Name                                 | Status |
+--------------------------------------+--------------------------------------+--------+
| 00acfc1f-2109-4e9c-b628-de7149b42dc3 | ubuntu-16.04-server-cloudimg-amd64 | Active |
| bf1abfb8-8bcc-4ce8-a9e8-3432b8ca546e | ubuntu1604_jenkins_node            | Active |
+--------------------------------------+--------------------------------------+--------+
```

4 Determine the UUID of the project.

```
openstack project list --domain default
```

The result lists the project UUID and name.

```
+----------------------------------+-------+
| ID                               | Name  |
+----------------------------------+-------+
| f33350f3844948fcb482ed6f5eef133d | admin |
+----------------------------------+-------+
```

5 Log in to the vSphere web client.

6 Go to vCenter and find the Project Folder with the UUID of the project.

7 In the Project Folder, find the template with the UUID of the image.

8 Right-click the template and select **Clone to Template** to open the Template to Template wizard:

a Enter a new name for the template.

b Choose a different host.

c Choose a datastore.

d Click **Finish** to complete the new template.

9 Right-click the original template and select **Delete from Disk**.

10 Right-click the template clone and select **Rename** to enter the original name as the new name.

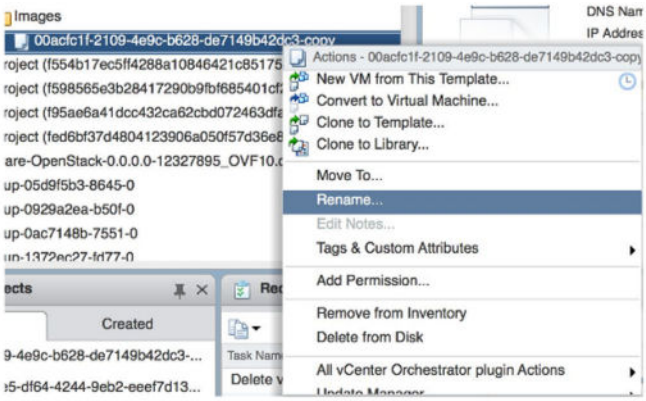

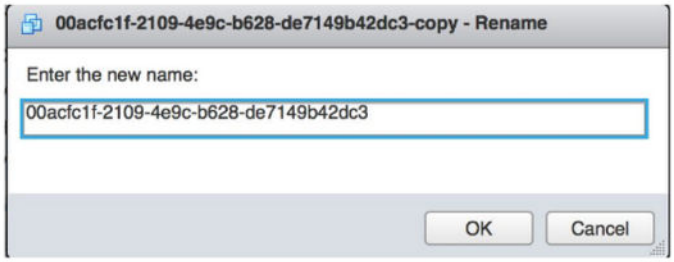

Configuring Images for Windows Guest Customization

You can configure images for Windows guest customization directly in the VMware Integrated OpenStack dashboard by applying the guest customization metadata to the Glance image used to create an instance.

The Windows guest customization feature provides an alternative to the cloudbase-init approach to enabling guest customization. If an image currently uses cloudbase-init, do not use the VMware Integrated OpenStack Windows guest customization feature.

Prerequisites

- Verify that you are logged in to the VMware Integrated OpenStack dashboard as a cloud administrator.
- Verify that you have an appropriate Windows OS image available in the Glance Image Service.
- Verify that the correct versions of the Microsoft System Preparation tools (sysprep) for each guest operating system you want to customize are installed in vSphere. See Installing the Microsoft Sysprep Tool in the vSphere product documentation.
- Verify that VMware Tools is installed on the source image.
- Verify that the image disk type property correctly reflects the image disk type prior to import.

  This applies only to images imported into Glance in VMware Integrated OpenStack versions earlier than 2.0. In version 2.0.x and later, image properties (such as disk type) are automatically introspected during the Glance import process.

Procedure

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

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

3 (Optional) Preview the Guest Customization Options metadata definition.

   a Select **Admin > System > Metadata Definitions**.

   b Click **Guest Customization Options**.

   c Click the **Contents** tab.

   You can only view the metadata definitions in the VMware Integrated OpenStack dashboard. You cannot modify the metadata.

4 Select **Admin > System > Images**.

5 Locate the Windows image to modify.

6 In the Actions column of the image listing, click the down arrow and select **Update Metadata**.

7 In the column under Available Metadata, expand the **Guest Customization Options** tab.

Note If the **Guest Customization Options** tab is not present, the related metadata properties might already be configured.

8 Click the plus sign (**+**) next to the guest customization option you want to add.

Tip You can add all the options simultaneously by the clicking the plus sign (**+**) on the top **Guest Customization Options** tab.

In the column under Existing Metadata, the newly added metadata properties appear.

Note You may need to scroll to the bottom of this column to see the newly added metadata properties.

9 Configure the metadata properties.

Metadata Property	Description
Auto logon count	Applies the `windows_logon_count` metadata property. Enter the number of times the machine can automatically logged in to as Administrator . Typically, this value is set to 1, but you can increase the value if your configuration requires multiple reboots. This value might be determined by the list of commands executed by the `GuiRunOnce` command.
Automatic logon	Applies the `windows_auto_logon` metadata property. If selected, the VM is automatically logged in to as Administrator.

Metadata Property	Description
Maximum number of connections	Applies the `windows_max_connect` metadata property. Enter the number of client licenses purchased for the Windows server being installed. **Note** This property is applied only if the `windows_license_mode` metadata property, described below, is set to `PerServer`.
Product Key	Applies the `windows_product_key` metadata property. Enter a valid serial number which is included in the answer file when mini-setup runs. **Note** This serial number is ignored if the original guest operating system was installed using a volume-licensed CD.
Server licensing mode	Applies the `windows_license_mode` metadata property. Select the licensing mode that matches your source image: `PerServer` or `PerSeat`.
Windows workgroup to join	Applies the `windows_join_workgroup` metadata property. Select the workgroup that the VM should join.

10 Click **Save**.

The image metadata is now configured for Windows guest customization and are applied for all future VMs that are created from this image.

Enable Live Resize

You can configure an image so that virtual machines deployed from the image can be resized while powered on.

You can enable live resize for the disk size, memory, vCPU, or any combination of these.

Note You cannot reboot a live resize-enabled instance with a volume attached. If you need to reboot the instance, detach the volume first.

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.

2 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.

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.

Table 9-1. Image Metadata in VMware Integrated OpenStack

Extra Spec	Description
`vmware_latency_sensitivity_level`	Specify the latency sensitivity level for virtual machines. Setting this key will adjust certain settings on virtual machines.
`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.
`quota_cpu_limit`	Specify the maximum CPU allocation in MHz. The value `0` indicates that CPU usage is not limited.
`quota_cpu_reservation`	Specifies the guaranteed CPU allocation in MHz.
`quota_cpu_reservation_percent`	Specifies 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`	Specifies 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`	Specifies the number of CPU shares allocated. If the `cpu_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`	Specifies 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`	Specifies the number of memory shares allocated. If the `memory_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`	Specifies 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`	Specifies the number of disk transaction shares allocated. If the `disk_io_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`	Specifies 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`	Specifies the number of virtual interface bandwidth shares allocated. If the `disk_io_shares_level` parameter is not set to `custom`, this value is ignored.

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 Block Storage.

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 Block Storage

It is a best practice to configure a backup service for the Block Storage (Cinder) component of OpenStack to prevent loss of data. You can configure Cinder to back up volumes to a network file system (NFS) server.

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

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

Prerequisites

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

Procedure

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

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

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

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

   c. Uncomment the `cinder_backup_share` parameter.

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

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

e If your NFS share does not match your VMware Integrated OpenStack deployment version, uncomment the `cinder_backup_mount_options` parameter and set it to your version of NFS.

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

4 Save the `custom.yml` file.

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

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

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

6 Verify that the backup service is operational

a Confirm that the Cinder backup service is running.

```
cinder service-list
```

b Create a test volume and back it up.

```
cinder create --display-name testvol
cinder backup-create --display-name testvol-backup testvol
```

c Check the NFS share to confirm that the backup file was created.

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*.

3 Recover OpenStack nodes by node or role.

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

```
sudo viocli recover {–n node... | –r role... [–n node...]}
```

Option	Description
–n	Recovers one or more specified nodes by name. 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.
–r	Recovers all nodes in one or more roles. You can specify `–n` in addition to this parameter to recover single nodes outside of the specified 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.

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

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

Option	Description
–n	Recovers one or more database nodes by name. You can specify `DB` nodes for HA deployments or the `ControlPlane` node for compact or tiny deployments.
–r	Recovers all database nodes. Specify `DB` for HA deployments or `ControlPlane` for compact or tiny deployments.
–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.

VMware Integrated 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/atop/atop_<YYYYMMDD>`	Atop is the OpenStack resource monitoring tool. Resource usage such as CPU, memory, and disk usage is sampled every 60 seconds and logs are stored in a directory with the date in YYYYMMDD format. Logs are rotated every 3 days by default.
`/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.

Name and Location	Description
/var/log/osvmw/osvmw.log	Logs osvmw service activity.
/var/log/viocli/viocli.log	Logs viocli (VMware Integrated OpenStack CLI) service activity.
/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.

OpenStack Controller Logs

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

Database Service Logs

Name and Location	Description
/var/log/syslog	General database logging including MySQL logging.
/var/log/rabbitmq/rabbit@database01.log	Logs general RabbitMQ database activity.

Name and Location	Description
/var/log/rabbitmq/shutdown_log	Logs RabbitMQ service shut-down activity.
/var/log/rabbitmq/startup_log	Logs RabbitMQ service start-up activity.

Compute and Loadbalancer Service Logs

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

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.

```
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
```

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.

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

3 Save the `custom.yml` file.

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

```
viocli deployment ––verbose 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 *
```

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

Troubleshoot Unable to Connect to Server

In the process of importing a vSphere VM into VMware Integrated OpenStack, connecting to the vAPI endpoint fails with an error.

Problem

Running DCLI to connect to the vAPI endpoint results in `ERROR: Unable to connect to the server`.

Cause

The vAPI service is not running.

Solution

1 Log into the OpenStack Manager Server as root.

2 Check the status of the vAPI service.

```
systemctl status vapi
```

The service appears inactive.

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

3 Restart the service.

```
systemctl restart vapi
```

4 Check the status of the vAPI service again.

```
systemctl status vapi
```

The service appears to restart.

```
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

With the service started, retry connecting to the VMware Integrated OpenStack vAPI endpoint. See Import Virtual Machines into VMware Integrated OpenStack.

VMware Integrated OpenStack APIs

12

VMware Integrated OpenStack includes several APIs that help you deploy and manage OpenStack.

This chapter includes the following topics:

Using the OpenStack Management Server APIs

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/login` and include `username=vcenter-user&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/401. If you have installed VMware Integrated OpenStack, you can also view the API specifications at `https://mgmt-server-ip:8443/swagger-ui.html`.

Using the Tenant Virtual Data Center vAPIs

VMware Integrated OpenStack includes vAPIs that you can use to manage tenant virtual data centers.

If you have logged in to the OpenStack Management Server, you can also manage tenant virtual data centers using the Data Center Command-Line Interface (DCLI) or the `viocli` utility. For information about the `viocli` utility, see viocli inventory-admin Command.

When using the vAPIs, you must authenticate with the vAPI endpoint using the administrator credentials for your vCenter Server instance.

You can use any HTTP client to send requests to the vAPI endpoint. This document uses cURL as an example.

Create a Tenant Virtual Data Center

```
curl -X POST -u vcserver-admin -H "Content-Type: application/json"
   https://mgmt-server-ip:9449/rest/vio/tenant/vdc
   -d '{
      "spec":{
         "compute":"compute-node",
         "display_name":"vdc-name",
         "project_id":"project-uuid",
         "cpu_limit":max-cpu-mhz,
         "cpu_reserve":min-cpu-mhz,
         "mem_limit":max-memory-mb,
         "mem_reserve":min-memory-mb
      }
   }'
```

The `cpu_limit`, `cpu_reserve`, `mem_limit`, and `mem_reserve` parameters are optional.

The ID of the new tenant virtual data center is returned in JSON format.

The equivalent DCLI command is as follows:

```
com vmware vio tenant vdc create --compute compute-node --display-name vdc-name --project-id project-
uuid [--cpu-limit max-cpu-mhz] [--cpu-reserve min-cpu-mhz] [--mem-limit max-memory-mb] [--mem-reserve
min-memory-mb]
```

Update a Tenant Virtual Data Center

```
curl -X PATCH -u vcserver-admin -H "Content-Type: application/json"
   https://mgmt-server-ip:9449/rest/vio/tenant/vdc/tenant-vdc-id
   -d '{
      "spec":{
         "compute":"compute01"
         "cpu_limit":max-cpu-mhz,
         "cpu_reserve":min-cpu-mhz,
         "mem_limit":max-memory-mb,
         "mem_reserve":min-memory-mb
      }
   }'
```

The `cpu_limit`, `cpu_reserve`, `mem_limit`, and `mem_reserve` parameters are optional.

The equivalent DCLI command is as follows:

```
com vmware vio tenant vdc update --compute compute-node --tvdc-id tenant-vdc-id [--cpu-limit max-cpu-
mhz] [--cpu-reserve min-cpu-mhz] [--mem-limit max-memory-mb] [--mem-reserve min-memory-mb]
```

List All Tenant Virtual Data Centers

```
curl –u vcserver–admin https://mgmt–server–ip:9449/rest/vio/tenant/vdc
```

The information is returned in JSON format.

The equivalent DCLI command is as follows:

```
com vmware vio tenant vdc list
```

Display Information About a Tenant Virtual Data Center

```
curl –u vcserver–admin https://mgmt–server–ip:9449/rest/vio/tenant/vdc/tenant–vdc–id
```

The status, provider ID, display name, and quotas of the tenant virtual data center are returned in JSON format.

The equivalent DCLI command is as follows:

```
com vmware vio tenant vdc get ––tvdc–id tenant–vdc–id
```

Delete a Tenant Virtual Data Center

```
curl –X POST –u vcserver–admin –H "Content–Type: application/json"
   https://mgmt–server–ip:9449/rest/vio/tenant/vdc/tenant–vdc–id?action=delete–tvdc
   –d '{
      "spec":{
         "compute":"compute–node"
      }
   }'
```

The `compute` parameter is optional. If you specify `compute`, the tenant virtual data center is deleted from the specified compute node only. If you do not specify `compute`, the tenant virtual data center is deleted from all compute nodes.

The equivalent DCLI command is as follows:

```
com vmware vio tenant vdc deletetvdc ––tvdc–id tenant–vdc–id [––compute compute–node]
```

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`.

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 [--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.
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`
`--verbose`	Optional	Displays output in verbose mode.

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.

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.
`--verbose`	Optional	Displays output in verbose mode.

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] [--verbose]

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] [--verbose]

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] [--verbose]

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] [--verbose]

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] [--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.
`--verbose`	Optional	Enter verbose mode.

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.
`--verbose`	Optional	Displays output in verbose mode.

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] [--verbose]`

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] [--verbose]`

Stops a deployment.

`viocli deployment pause [-d NAME] [-p] [--verbose]`

Pauses a deployment.

`viocli deployment resume [-d NAME] [-p] [--verbose]`

Resumes a paused deployment.

`viocli deployment reset_status [-d NAME] [-p] [--verbose]`

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] [--verbose]`

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] [--verbose]

Updates the post-deployment configuration.

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

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] [--verbose]

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* or --organization_unit_name *ORG-UNIT*	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] [--verbose]

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*] [-nrl] [--recent-logs] [-p] [--verbose]

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` ■ `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] [--verbose]

Returns the name of the default deployment.

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

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.
`--verbose`	Optional	Displays output in verbose mode.

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

viocli enable-tvd Command

Use the `viocli enable-tvd` command to add NSX-T networking support to a VMware Integrated OpenStack deployment that was previously deployed with NSX-V networking.

Important This command will update your `custom.yml` file or automatically generate a `custom.yml` file if the file does not exist in your environment. After running the `viocli enable-tvd` command, do not delete or overwrite `custom.yml` or your configuration will be discarded.

The `viocli enable-tvd` command uses the following syntax.

```
viocli enable-tvd [-d NAME] --nsx-mgr MANAGER-IP --nsx-user USERNAME --nsx-passwd PASSWORD [--nsx-
insecure {true | false}] [--nsx-ca-file CA-FILE] [--nsx-overlay-tz OVERLAY-TZ] [--nsx-vlan-tz VLAN-TZ]
[--nsx-tier0-rt TIER0-ROUTER] [--nsx-dhcp-profile DHCP-PROFILE] [--nsx-md-proxy MD-PROXY] [--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.
`--nsx-mgr` *MANAGER-IP*	Mandatory	IP address of the NSX Manager of your NSX-T deployment.
`--nsx-user` *USERNAME*	Mandatory	User name of the NSX Manager administrator.
`--nsx-passwd` *PASSWORD*	Mandatory	Password for the NSX Manager administrator.
`--nsx-insecure {true | false}`	Optional	Specifies whether to verify the certificate of the NSX Manager server. If you do not include this option, `true` is used by default.
`--nsx-ca-file` *-CA-FILE*	Optional	CA bundle files to use in verifying the certificate of the NSX Manager server. This option is ignored if you include the `--nsx-insecure true` option.
`--nsx-overlay-tz` *OVERLAY-TZ*	Optional	Name or UUID of the default NSX-T overlay transport zone used for creating tunneled isolated Neutron networks. You must create the zone in NSX-T before using the plugin.
`--nsx-vlan-tz` *VLAN-TZ*	Optional	Name or UUID of the default NSX-T VLAN transport zone used for bridging between Neutron networks if no physical network has been specified.
`--nsx-tier0-rt` *TIER0-ROUTER*	Optional	Name or UUID of the default tier-0 router used to connect to tier-1 logical routers and configure external networks.
`--nsx-dhcp-profile` *DHCP-PROFILE*	Optional	Name or UUID of the NSX-T DHCP profile used to enable native DHCP service. You must create the profile in NSX-T before using the plugin.

Parameter	Mandatory or Optional	Description
`--nsx-md-proxy` *MD-PROXY*	Optional	Name or UUID of the NSX-T metadata proxy server used to enable native metadata service. You must create the proxy server in NSX-T before using the plugin.
`--verbose`	Optional	Displays output in verbose mode.

Note If you do not include the `--nsx-ca-file`, `--nsx-overlay-tz`, `--nsx-vlan-tz`, `--nsx-tier0-rt`, `--nsx-dhcp-profile`, or `--nsx-md-proxy` parameters, the system will attempt to determine the correct information automatically. If the command fails, retry with the parameters included.

You can also run `viocli enable-tvd -h` or `viocli enable-tvd --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 current deployment is used by default.
`--verbose`	Optional	Displays output in verbose mode.

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* `[--verbose]`

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	Specifies the local path or URL to the agent installer package.
`-c` *PROP-FILE* or `--config` *PROP-FILE*	Mandatory	Specifies the local path to the agent configuration file.

`viocli epops uninstall [-d NAME] [--verbose]`

Uninstalls the End Point Operations Management agent.

`viocli epops reconfig [-d NAME] -c PROP-FILE [--verbose]`

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	Specifies the local path to the agent configuration file.

`viocli epops start [-d NAME] [--verbose]`

Starts the End Point Operations Management agent.

`viocli epops stop [-d NAME] [--verbose]`

Stops the End Point Operations Management agent.

viocli federation Command

Use the `viocli federation` command to configure Keystone federated identity in your VMware Integrated OpenStack deployment.

The `viocli federation` command can perform various actions on identity provider (IdP) metadata, Keystone service providers, and Keystone identity providers. The following parameters apply to all actions on all components.

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.
`--verbose`	Optional	Displays output in verbose mode.

You can run `viocli federation -h` or `viocli federation --help` to display the parameters for the command. You can also use the `-h` or `--help` option on any component or action to display relevant parameters. For example, `viocli federation idp-metadata -h` will show parameters for the `idp-metadata` component and `viocli federation idp-metadata set -h` will show parameters for the `set` action on that component.

Identity Provider Metadata

The actions that `viocli federation` supports for identity provider metadata are listed as follows.

`viocli federation idp-metadata clear [-d NAME] [-p] [--verbose]`

Removes identity provider metadata.

`viocli federation idp-metadata set [-d NAME] [-p] [--verbose]`

Sets updated identity provider metadata. You are prompted to enter information for the organization and contact person.

`viocli federation idp-metadata show [-d NAME] [--json JSON | --pretty PRETTY] [-p] [--verbose]`

Displays identity provider metadata. The following additional parameters apply to the `show` 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.

Note After updating or removing metadata, you must run the `viocli identity configure` command to make your changes take effect.

Keystone Service Providers

The actions that `viocli federation` supports for Keystone service providers are listed as follows.

viocli federation service-provider add [-d *NAME*] [--type *SP-TYPE*] [-p] [--verbose]

Adds a service provider. You are prompted to enter Keystone information. The following additional parameters apply to the `add` action.

Parameter	Mandatory or Optional	Description
`--type` *SP-TYPE*	Optional	Specifies the type of service provider to add.

viocli federation service-provider remove [-d *NAME*] --id *SP-ID* [-p] [--verbose]

Removes a service provider. The following additional parameters apply to the `remove` action.

Parameter	Mandatory or Optional	Description
`--id` *SP-ID*	Mandatory	Identifier of the service provider to remove. You can run the `viocli federation service-provider list` command to display the identifier of each service provider.

viocli federation service-provider edit [-d *NAME*] --id *SP-ID* [-p] [--verbose]

Modifies the configuration of a service provider. The following additional parameters apply to the `edit` action.

Parameter	Mandatory or Optional	Description
`--id` *SP-ID*	Mandatory	Identifier of the service provider to modify. You can run the `viocli federation service-provider list` command to display the identifier of each service provider.

viocli federation service-provider list [-d *NAME*] [--json JSON | --pretty PRETTY] [-p] [--verbose]

Displays information about all service providers. 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 federation service-provider show [-d NAME] --id SP-ID [--json JSON | --pretty PRETTY] [-p] [--verbose]`

Displays detailed information about a service provider. The following additional parameters apply to the `show` action.

Parameter	Mandatory or Optional	Description
`--id` *SP-ID*	Mandatory	Identifier of the service provider.
`--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.

Keystone Identity Providers

The actions that `viocli federation` supports for Keystone identity providers are listed as follows.

`viocli federation identity-provider add [-d NAME] [--type {keystone | saml2 | vidm}] [-p] [--verbose]`

Adds a service provider. You are prompted to enter Keystone information. The following additional parameters apply to the `add` action.

Parameter	Mandatory or Optional	Description
`--type {keystone | saml2 | vidm}`	Optional	Specifies the type of identity provider to add. If you do not include this option in the command, you will be prompted to enter a value.

`viocli federation identity-provider remove [-d NAME] --id IDP-ID [-p] [--verbose]`

Removes an identity provider. The following additional parameters apply to the `remove` action.

Parameter	Mandatory or Optional	Description
`--id` *IDP-ID*	Mandatory	Identifier of the identity provider to remove. You can run the `viocli federation identity-provider list` command to display the identifier of each identity provider.

`viocli federation identity-provider edit [-d` *NAME*`] --id` *IDP-ID* `[-p] [--verbose]`

Modifies the configuration of an identity provider. The following additional parameters apply to the `edit` action.

Parameter	Mandatory or Optional	Description
`--id` *IDP-ID*	Mandatory	Identifier of the identity provider to modify. You can run the `viocli federation identity-provider list` command to display the identifier of each identity provider.

`viocli federation identity-provider list [-d` *NAME*`] [--json JSON | --pretty PRETTY] [-p] [--verbose]`

Displays information about all identity providers. 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 federation identity-provider show [-d` *NAME*`] --id` *IDP-ID* `[--json JSON | --pretty PRETTY] [-p] [--verbose]`

Displays detailed information about an identity provider. The following additional parameters apply to the `show` action.

Parameter	Mandatory or Optional	Description
`--id` *IDP-ID*	Mandatory	Identifier of the identity provider.
`--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 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.
`--verbose`	Optional	Displays output in verbose mode.

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] [--verbose]

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] [--verbose]

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 configure [–d *NAME*] [–p] [––verbose]

Configures identity sources for your deployment.

viocli identity edit [–d *NAME*] ––id *DOMAIN* [–p] [––verbose]

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] [––verbose]

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] [––verbose]

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.
`--verbose`	Optional	Displays output in verbose mode.

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] [--verbose]
```

Lists orphaned Nova instances.

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

Lists orphaned vSphere virtual machines.

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

Lists orphaned shadow virtual machines.

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

Removes orphaned vSphere virtual machines.

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

Removes orphaned vSphere virtual machines.

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

Removes orphaned shadow virtual machines.

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

Lists hypervisors with detailed information.

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

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] [--verbose]

Synchronizes the availability zones in the environment with the specified map. The following additional parameters are supported.

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*] [--cpu-limit *CPU-LIMIT*] [--cpu-shares *CPU-SHARES*] [--mem-reserve *MEMORY*] [--mem-limit *MEM-LIMIT*] [--mem-shares *MEM-SHARES*] [--json | --pretty] [--all] [--force] [--no-grace-period] [--verbose]

Create a tenant virtual data center (VDC) with the specified settings. The following additional parameters are supported.

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*	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-LIMIT*	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.

Parameter	Mandatory or Optional	Description
`--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] [--verbose]

Lists tenant VDCs.

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

Displays detailed information about the specified tenant VDC. The following additional parameters are supported.

Parameter	Mandatory or Optional	Description
`--id` *ID*	Mandatory	Identifier of a tenant VDC.

viocli inventory-admin delete-tenant-vdc [-d *NAME*] --id *ID* [--compute *COMPUTE-NODE*] [--json | --pretty] [--all] [--force] [--no-grace-period] [--verbose]

Deletes the specified tenant VDC. The following additional parameters are supported.

Parameter	Mandatory or Optional	Description
`--id` *ID*	Mandatory	Identifier of a tenant VDC.
`--compute` *COMPUTE-NODE*	Optional	Compute node from which to delete the VDC. If you do not enter a value, the VDC is deleted from all compute nodes.

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

Updates the configuration of the specified tenant VDC. The following additional parameters are supported.

Parameter	Mandatory or Optional	Description
`--compute` *COMPUTE-NODE*	Mandatory	Compute node on which the VDC is running.
`--id` *VDC-ID*	Mandatory	Identifier of the tenant VDC.
`--cpu-reserve` *CPU-MIN*	Optional	CPU cycles in MHz to reserve for the VDC.
`--cpu-limit` *CPU-MAX*	Optional	Maximum limit for CPU usage on the VDC (in MHz). The value `-1` indicates that CPU usage is not limited.
`--mem-reserve` *MEMORY-MIN*	Optional	Memory in megabytes to reserve for the VDC.
`--mem-limit` *MEMORY-MAX*	Optional	Maximum limit for memory consumption on the VDC (in megabytes). The value `-1` indicates that memory usage is not limited.

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] [--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.
`-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 command. 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`
`-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
`--verbose`	Optional	Displays output in verbose mode.

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 [--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.
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`
`--verbose`	Optional	Displays output in verbose mode.

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} [--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.
`--verbose`	Optional	Displays output in verbose mode.

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]
```

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.

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.
`--verbose`	Optional	Displays output in verbose mode.

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] [––verbose]

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] [––verbose]

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] [––verbose]

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] [--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.
`--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.
DEST-DS-NAME	Mandatory	Specifies the destination datastore.

Parameter	Mandatory or Optional	Description
`--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.
`--verbose`	Optional	Displays output in verbose mode.

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 [--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.
`-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.
`--verbose`	Optional	Displays output in verbose mode.

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] [––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.
`–p` or `––progress`	Optional	Shows the progress of the current operation.
`––verbose`	Optional	Displays output in verbose mode.

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.