

# VMware Integrated OpenStack User Guide

VMware Integrated OpenStack 4.1



vmware®

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

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

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

[docfeedback@vmware.com](mailto:docfeedback@vmware.com)

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

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

# Contents

- About This Book 5
  
- 1 Log In to the VMware Integrated OpenStack Dashboard 6**
  
- 2 Managing Images for the Image Service 8**
  - Import Images Using the Horizon Dashboard 8
  - Import Images in Supported Formats Using the CLI 9
  - Modify Image Settings 11
  - Delete an Existing Image 11
  
- 3 Configuring Access and Security for Instances 12**
  - Working with Security Groups 12
  - Working with Key Pairs 15
  - Allocate a Floating IP to an Instance 16
  
- 4 Working with Networks 18**
  - Create a Network 18
  - Create a Router 19
  
- 5 Working with Instances in OpenStack 21**
  - Start an OpenStack Instance from an Image 21
  - Start an OpenStack Instance from a Snapshot 23
  - Connect to an Instance by Using SSH 24
  - Track Instance Use 24
  - Create a Snapshot from an Instance 24
  - Using Affinity and Anti-Affinity to Place OpenStack Instances 25
  
- 6 Working with Volumes 28**
  - Create a Volume 28
  - Modify Existing Volumes 29
  - Delete Existing Volumes 29
  - Attach a Volume to an Instance 30
  - Detach a Volume 30
  - Create a Snapshot from a Volume 31
  
- 7 Working with Orchestration and Stacks 32**
  - Start a New Orchestration Stack 32
  - Modify an Orchestration Stack 34

[Delete an Orchestration Stack](#) 34

# About This Book

*VMware Integrated OpenStack User Guide* shows you how to perform VMware Integrated OpenStack cloud end-user tasks in VMware Integrated OpenStack, including how to create and manage instances, volumes, snapshots, images, and networks.

As a VMware Integrated OpenStack cloud end user, you can provision your own resources within the limits that administrators set.

## Intended Audience

This guide is for cloud users who want to create and manage resources with an OpenStack deployment that is fully integrated with VMware<sup>®</sup> vSphere<sup>®</sup>. To do so successfully, verify that you are 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>.

# Log In to the VMware Integrated OpenStack Dashboard

1

You access the user and administrative controls for your VMware Integrated OpenStack deployment through the VMware Integrated OpenStack dashboard. The dashboard enables you to create and manage instances, images, user accounts, and volumes, among other tasks.

To log in to the dashboard, you must obtain the host name or IP address for the VMware Integrated OpenStack dashboard from your OpenStack operator. This is the public virtual IP created when deploying up the VMware Integrated OpenStack in vSphere.

## Prerequisites

- Verify that you have a user account that was set up by an administrative user.
- Verify that you have a browser with JavaScript and cookies enabled.

## Procedure

- 1 In a browser window, navigate to the host name or IP address for the VMware Integrated OpenStack dashboard.

A certificate warning might appear the first time you access the URL. To bypass the warning, verify the certificate or add an exception.

- 2 On the Log In page, enter the domain name, your user name and password.
- 3 Click **Sign In**.

You are now logged in. The Project tab appears, opened to the default Overview page.

Figure 1-1. VMware Integrated OpenStack Overview Page

The screenshot displays the VMware Integrated OpenStack Overview Page. The interface includes a top navigation bar with the VMware logo, a 'service' dropdown, a user profile 'writer\_andy', and a 'Sign Out' link. A left sidebar shows navigation options: Project, Compute (selected), Overview, Instances, Volumes, Images, Access & Security, and Network. The main content area is titled 'Overview' and features a 'Limit Summary' section with seven circular progress indicators showing usage against limits: Instances (0 of 10), VCPUs (0 of 20), RAM (0 Bytes of 50.0GB), Floating IPs (0 of 50), Security Groups (1 of 10), Volumes (0 of 10), and Volume Storage (0 Bytes of 1000.0GB). Below this is a 'Usage Summary' section with a date range selector (From: 2014-12-01, To: 2014-12-22) and a 'Submit' button. The usage summary shows 0 Active Instances, 0 Active RAM, 0 Bytes, 0.00 VCPU-Hours, and 0.00 GB-Hours. At the bottom is a 'Usage' table with columns for Instance Name, VCPUs, Disk, RAM, and Uptime, which is currently empty.

# Managing Images for the Image Service

# 2

In the OpenStack context, an image is a file that contains a virtual disk from which you can install an operating system on a VM. You create an instance in your OpenStack cloud by using one of the images available. The VMware Integrated OpenStack Image Service component natively supports images that are packaged in the ISO, OVA, and VMDK formats.

If you have existing images in vSphere that you want to use in OpenStack, you can export them in one of the supported formats and upload them to the Image Service. If you obtain an image that is in an unsupported format, you can convert it as part of the import process. Unsupported formats are RAW, QCOW2, VDI, and VHD.

This chapter includes the following topics:

- [Import Images Using the Horizon Dashboard](#)
- [Import Images in Supported Formats Using the CLI](#)
- [Modify Image Settings](#)
- [Delete an Existing Image](#)

## Import Images Using the Horizon Dashboard

You can import images directly in the VMware Integrated OpenStack Horizon dashboard.

### Prerequisites

- Verify that the image is packaged in the ISO, VMDK, OVA, RAW, QCOW2, VDI, or VHD format.
- If the source image format is RAW, QCOW2, VDI, or VHD, verify that the source image is hosted on a server without credentials to allow plain HTTP requests.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Images**.
- 4 On the Images page, click **Create Image**.



## 5 Configure the image.

| Option                   | Action   |
|--------------------------|--|
| <b>Name</b>              | Enter a name for the new image.  |
| <b>Description</b>       | (Optional) Enter a description for the new image.  |
| <b>Image Source</b>      | Select the image source.<br>If the source image format is RAW, QCOW2, VDI, or VHD, you must select the Image Location option.  |
| <b>Disk Format</b>       | Select the disk format.  |
| <b>Disk Type</b>         | Select the disk type.<br>Images in the RAW, QCOW2, VDI, and VHD formats are automatically introspected to capture their properties and converted to the VMDK format during the import process. |
| <b>Adapter Type</b>      | Select the adapter type.   |
| <b>Architecture</b>      | Accept the default.  |
| <b>OS Type</b>           | Select the type of operating system.   |
| <b>Minimum Disk (GB)</b> | Specify the minimum disk size for the image in GB.   |
| <b>Minimum RAM (GB)</b>  | Specify the minimum RAM for the image.   |
| <b>Public</b>            | Select to make the image visible and available to all tenants.   |
| <b>Protected</b>         | Select to prevent the image from being deleted.  |

## 6 Click **Create Image**.

The Images page now includes the newly added image.

The image is now ready for deployment in OpenStack instances.

# Import Images in Supported Formats Using the CLI

You can make images available for use in instances by importing images to the Image Service datastore .

To import an image in a non-supported format such as RAW, QCOW2, VDI, or VHD, see the *VMware Integrated OpenStack Administrator Guide*.

### Prerequisites

- Verify that you configured one or more Image Service datastores.
- Obtain the image, for example, `ubuntuLTS-sparse.vmdk`.
- Verify that the images are packaged in the ISO, VMDK, or OVA format.

### Procedure

- 1 Log in to the OpenStack management cluster as a user with administrative privileges to upload the image to the Image Service component.

## 2 Run the `glance image-create` command to obtain, define, and import the image.

```
glance --os-auth-token $token --os-image-url http://123.456.7.8:9292 \
  image-create name="ubuntu-sparse" \
  disk_format=vmdk \
  container_format=bare \
  --visibility="public" \
  --property vmware_adapertype="lsiLogicsas" \
  --property vmware_disktype="sparse" \
  --property vmware_ostype="ubuntu64Guest" < ubuntuLTS-sparse.vmdk
```

This example uses the following parameters and settings.

| Parameter or Setting  | Description  |               |  |                     |  |                        |   |
|---|--|---------------|--|---------------------|--|------------------------|---|
| <code>--os-image-url</code><br><code>http://123.456.7.8:9292</code>     | The URL of the source image.   |               |  |                     |  |                        |   |
| <code>name="ubuntu-sparse"</code>                                       | The name of the source image, in this case, <b>ubuntu-sparse</b> .   |               |  |                     |  |                        |   |
| <code>disk_format=vmdk</code>   | The disk format of the source image. You can specify ISO, VMDK, or OVA.  |               |  |                     |  |                        |   |
| <code>container_format=bare</code>                                      | The container format indicates if the image is in a format that contains metadata about the actual virtual machine. Because the container format string is not currently used by Glance, it is recommended to specify <b>bare</b> for this parameter.  |               |  |                     |  |                        |   |
| <code>--visibility="public"</code>                                      | The privacy setting for the image in OpenStack. When set to <b>public</b> , the image is available to all users. When set to <b>private</b> , the image is available only to the current user.   |               |  |                     |  |                        |   |
| <code>--property</code><br><code>vmware_adapertype="lsiLogicsas"</code> | <p>During import, the VMDK disk is introspected to capture its adapter type property.</p> <p>You also have the option of using the <code>vmware_adapertype</code> to specify adapter type.</p> <p><b>Note</b> If you are using a disk with the paraVirtual or LSI Logic SAS adapter type, it is recommend that you use this parameter. For example, <code>vmware_adapertype= lsiLogicsas</code> or <code>vmware_adapertype= paraVirtual</code>.</p>  |               |  |                     |  |                        |   |
| <code>--property vmware_disktype="sparse"</code>                        | <p>During import, the VMDK disk type is introspected to capture its disk type property.</p> <p>You also have the option of specifying disk type using the <code>vmware_disktype</code> property.</p> <table border="0"> <tr> <td><b>sparse</b></td> <td>This disktype property applies to monolithic sparse disks.</td> </tr> <tr> <td><b>preallocated</b></td> <td>This disktype property applies to VMFS flat disks, including thick, zeroedthick, or eagerzeroedthick. This is the default property if none is specified.</td> </tr> <tr> <td><b>streamOptimized</b></td> <td>This disktype property applies to Monolithic Sparse disks, optimized for streaming. You can convert disks dynamically to and from this format with minimal computational costs.</td> </tr> </table> | <b>sparse</b> | This disktype property applies to monolithic sparse disks. | <b>preallocated</b> | This disktype property applies to VMFS flat disks, including thick, zeroedthick, or eagerzeroedthick. This is the default property if none is specified. | <b>streamOptimized</b> | This disktype property applies to Monolithic Sparse disks, optimized for streaming. You can convert disks dynamically to and from this format with minimal computational costs. |
| <b>sparse</b>   | This disktype property applies to monolithic sparse disks.   |               |  |                     |  |                        |   |
| <b>preallocated</b>   | This disktype property applies to VMFS flat disks, including thick, zeroedthick, or eagerzeroedthick. This is the default property if none is specified.   |               |  |                     |  |                        |   |
| <b>streamOptimized</b>  | This disktype property applies to Monolithic Sparse disks, optimized for streaming. You can convert disks dynamically to and from this format with minimal computational costs.  |               |  |                     |  |                        |   |

| Parameter or Setting                                  | Description   |
|---|---|
| <code>--property vmware_ostype="ubuntu64Guest"</code> | The name of the image file after it is imported to the Image Service. In the example above, the resulting name will be <code>ubuntuLTS-sparse.vmdk</code> . |

- 3 (Optional) In the Compute component, confirm that the image was successfully imported.

```
$ glance image-list
```

The command returns a list of all images that are available in the Image Service.

## Modify Image Settings

After an image is loaded, you can modify the image settings, such as image name, description, and the public and protected settings.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Images**.
- 4 Select the image to edit.
- 5 In the Actions column, click **Edit Images**.
- 6 Modify the settings as necessary.
- 7 Click **Update Image**.

The Images page redisplay with the changed information.

## Delete an Existing Image

Deleting an image is permanent and cannot be reversed. You must have administrative permissions to delete an image.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Images**.
- 4 Select one or more images to delete.
- 5 Click **Delete Images**.
- 6 Confirm the deletion at the prompt.

# Configuring Access and Security for Instances

# 3

Before you start instances, configure access and security settings. For example, SSH access and ICMP access are not enabled by default.

|                        |   |
|------------------------|---|
| <b>Security groups</b> | Enable users to ping and use SSH to connect to the instance. Security groups are sets of IP filter rules that define networking access and are applied to all instances in a project.   |
| <b>Key pairs</b>       | SSH credentials that are injected into an instance when it starts. To use key pair injection, the image that the instance is based on must contain the cloud-init package. Each project must have at least one key pair. If you generated a key pair with an external tool, you can import it into OpenStack. You can use the key pair for multiple instances that belong to a project. |
| <b>Floating IPs</b>    | When you create an instance in OpenStack, it is assigned a fixed IP address in the network. This IP address is permanently associated with the instance until the instance is terminated. You can also attach to an instance a floating IP address whose association can be modified.   |

This chapter includes the following topics:

- [Working with Security Groups](#)
- [Working with Key Pairs](#)
- [Allocate a Floating IP to an Instance](#)

## Working with Security Groups

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

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

---

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

---

**Important** For deployments with VMware NSX-T, the maximum number of security groups per port is 9.

---

## About the Default Security Group

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

## Create a Security Group

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

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

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Security Groups** tab.
- 5 Click **Create Security Group**.
- 6 Enter a name and description for the new group, and click **Create Security Group**.

The new group appears in the list on the **Security Group** tab.

- 7 Configure the rules for the new group.
  - a Select the new security group and click **Manage Rules**.
  - b Click **Add Rule**.
  - c From the **Rule** drop-down menu, select the rule to add.

The subsequent fields might change depending on the rule you select.
  - d If applicable, specify **Ingress** or **Egress** from the **Direction** drop-down menu.
  - e After you complete the rule definition, click **Add**.
- 8 Configure additional rules if necessary.
- 9 Click the **Access & Security** tab to return to the main page.

## Modify the Rules for an Existing Security Group

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

**Procedure**

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Security Groups** tab.
- 5 Select the security group to modify and click **Manage Rules**.
- 6 To remove a rule, select the rule and click **Delete Rule**.
- 7 To add a rule, click **Add Rule** and select the custom rule to add from the **Rule** drop-down menu.

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

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

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

- c Select the kind of access to allow.

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

- 8 Click **Add**.

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

## Enabling SSH and ICMP Access

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

**Procedure**

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.

- 4 Click the **Security Groups** tab, select the default security group, and click **Manage Rules**.
- 5 Click **Add Rule** and configure the rules to allow SSH access.

| Control | Value     |
|---------|-----------|
| Rule    | SSH       |
| Remote  | CIDR      |
| CIDR    | 0.0.0.0/0 |

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

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

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

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

- 8 Click **Add**.

Instances will now accept all incoming ICMP packets.

## Working with Key Pairs

Key pairs are SSH credentials that are injected into an instance when it starts.

To use key pair injection, the image that the instance is based on must contain the cloud-init package. Each project should have at least one key pair. If you generated a key pair with an external tool, you can import it into OpenStack. You can use the key pair for multiple instances that belong to a project.

### Add a Key Pair

Key pairs are SSH credentials that are injected into an instance when it starts. You can create or import key pairs.

You must provide at least one key pair for each project.

#### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.

- 4 Click the **Key Pairs** tab, which lists the key pairs available for the current project.
- 5 Click **Create Key Pair**.
- 6 Enter a name for the new key pair, and click **Create Key Pair**.
- 7 Download the new key pair at the prompt.
- 8 On the main **Key Pairs** tab, confirm that the new key pair is listed.

## Import a Key Pair

Key pairs are SSH credentials that are injected into an instance when it starts. You can create or import key pairs.

You must provide at least one key pair for each project.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Key Pairs** tab, which lists the key pairs available for the current project.
- 5 Click **Import Key Pair**.
- 6 Enter the name of the key pair. .
- 7 Copy the public key to the Public Key text box and click **Import Key Pair**.
- 8 Return to the main **Key Pairs** tab to confirm that the imported key pair is listed.

## Allocate a Floating IP to an Instance

You can attach a floating IP address to an instance in addition to the fixed IP address that is assigned when it is created. Unlike fixed IP addresses, you can modify floating IP address associations at any time, regardless of the state of the instances involved.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Access & Security**.
- 4 Click the **Floating IPs** tab, and click **Allocate IP to Project**.
- 5 Choose the pool from which to pick the IP address and click **Allocate IP**.



- Click **Associate** in the Floating IPs list and configure the floating IP associations settings.

| Option                 | Description  |
|------------------------|--|
| IP Address             | Click the plus sign to add an IP address.  |
| Ports to be associated | Select a port from the list. The list shows all the instances with their fixed IP addresses. |

- Click **Associate**.
- (Optional) To disassociate a floating IP address from an instance, click the **Floating IPs** tab, and click **Disassociate** in the Actions column for the IP address. .
- To release the floating IP address back into the pool of addresses, click **More** and select **Release Floating IP**.
- Click the **Floating IPs** tab and select the IP address.
- Click **Release Floating IPs**.

# Working with Networks

The OpenStack Networking service provides a scalable system for managing the network connectivity in an OpenStack cloud deployment. It can react to changing network needs, for example, creating and assigning new IP addresses. You can also configure logical routers to connect the different networks within your VMware Integrated OpenStack deployment.

For more information about how to manage networks, see the *VMware Integrated OpenStack Administrator Guide*.

This chapter includes the following topics:

- [Create a Network](#)
- [Create a Router](#)

## Create a Network

The OpenStack Networking service component is a scalable system for managing network connectivity within your VMware Integrated OpenStack deployment. With the VMware Integrated OpenStack dashboard, you can quickly create logical networks.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Network > Networks**.  
The Networks page lists the networks that are currently configured.
- 4 Click **Create Network**.
- 5 On the **Network** tab, enter a name for the new network.
- 6 (Optional) Select **Admin State** to have the network forward packets.
- 7 Click **Next**.

## 8 Configure the subnet.

| Option                 | Action   |
|------------------------|--|
| <b>Create Subnet</b>   | Select to create a subnet. You do not have to specify a subnet when you create a network, but if you do not, attached instances receive an Error status. To create a network without a subnet, deselect <b>Create Subnet</b> . |
| <b>Subnet Name</b>     | (Optional) Enter a name for the subnet.  |
| <b>Network Address</b> | If you create a subnet associated with the new network, specify the IP address for the subnet using the CIDR format, for example, 192.168.0.0/24.  |
| <b>IP Version</b>      | Select IPv4 or IPv6 from the drop-down menu.   |
| <b>Gateway IP</b>      | Enter the IP address for a specific gateway.   |
| <b>Disable Gateway</b> | (Optional) Select to disable a gateway IP address.   |

9 Click **Next** to access the settings on the **Subnet Detail** tab.

10 (Optional) if you selected the Create Subnet option on the previous tab, enter the subnet settings.

| Option                  | Description  |
|-------------------------|--|
| <b>Enable DHCP</b>      | (Optional) Select this option to enable DHCP. Consult with your network administrator. |
| <b>Allocation Pools</b> | Specify IP address pools for use by devices in the new network.                        |
| <b>DNS Name Servers</b> | Specify DNS servers for the new network.   |
| <b>Host Routes</b>      | Specify the IP address for the host routes.  |

11 Click **Create**.

When you start a new instance, this network will be available.

## Create a Router

With the VMware Integrated OpenStack dashboard, you can create logical routers. You use logical routers to connect the networks in your OpenStack deployment.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Network > Routers**.

The Routers page lists the routers that are currently configured.

- 4 Click **Create Router**.
- 5 Provide a name for the router and click **Create Router**.

The new router appears in the list on the Routers page. You can now complete the router configuration.

- 6 Click **Set Gateway** in the Actions column of the new router.

- 7 Select a network from the drop-down menu, and click **Set Gateway**.

The Router Name and Router ID text boxes are automatically populated.

- 8 Connect the router to a private network.

- a Click the router name on the Routers page.

- b Click **Add Interface**.

- c Select a subnet from the drop-down menu.

- d (Optional) Enter the router interface IP address for the selected subnet.

If you do not set this value, the first host IP address in the subnet is used by default.

- e Click **Add Interface**.

You successfully created the router. You can view the new topology on the Network Topology page.

# Working with Instances in OpenStack

# 5

Instances are virtual machines that run in the cloud.

You can start an instance from the following sources:

- Images uploaded to the OpenStack Image Service. See [Chapter 2 Managing Images for the Image Service](#).
- An image that you copied to a persistent volume. The instance starts from the volume, which the cinder-volume API provides through iSCSI. See [Attach a Volume to an Instance](#).

This chapter includes the following topics:

- [Start an OpenStack Instance from an Image](#)
- [Start an OpenStack Instance from a Snapshot](#)
- [Connect to an Instance by Using SSH](#)
- [Track Instance Use](#)
- [Create a Snapshot from an Instance](#)
- [Using Affinity and Anti-Affinity to Place OpenStack Instances](#)

## Start an OpenStack Instance from an Image

When you start an instance from an image, OpenStack creates a local copy of the image on the compute node where the instance starts. You can observe OpenStack instances in vSphere as VMs, but you must manage them in OpenStack.

### Prerequisites

Verify that images, flavors, block storage, and networks are configured and available to start an instance.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Images**.

The Images page lists the images available to the current user.

4 In the Actions column of the image, click **Launch**.

5 On the **Details** tab .

| Setting              | Description  |
|----------------------|--|
| Availability Zone    | Set by default to the availability zone that the cloud provider gives, for example: <b>nova</b> .  |
| Instance Name        | Name assigned to the VM. This value is a label and is not validated. When you create an instance, a UUID is assigned to the instance. When you view the VM in vSphere, you can identify the instance by the UUID but not by the instance name. |
| Flavor               | Size of the instance to start. The cloud administrator defines and manages flavors.  |
| Instance Count       | Number of instances started. The default is <b>1</b> .   |
| Instance Boot Source | Select <b>Boot from image</b> , and select the image from the list.  |

6 On the **Access & Security** tab of the Launch Instance dialog box .

| Setting         | Description   |
|-----------------|---|
| Key Pair        | Specify a key pair.<br>If the image uses a static root password or a static key set, you do not need to provide a key pair to start the instance, but using a key pair is a best practice.  |
| Security Groups | Select the security groups to be assigned to the instance.<br>Security groups are sets of rules that determine which incoming network traffic is forwarded to instances. If you did not create security groups, you can assign only the default security group to the instance. |

7 On the **Networking** tab, click the **+** icon in the Available Networks field to add a network to the instance.

8 (Optional) On the **Post-Creation** tab, specify a customization script that runs after the instance launches.

9 On the **Advanced Options** tab, select the type of disk partition from the drop-down list.

| Setting   | Description  |
|-----------|--|
| Automatic | The entire disk is a single partition and resizes.           |
| Manual    | Enables faster build times but requires manual partitioning. |

10 Click **Launch**.

The new instance starts on a node in the Compute cluster.

11 To view the new instance, select **Project > Compute > Instances**.

The Instances page shows the instance name, its private and public IP addresses, size, status, task, and power state.

## Start an OpenStack Instance from a Snapshot

You can start an instance from an instance snapshot. You can observe OpenStack instances in vSphere as VMs, but you can only manage them in OpenStack.

### Prerequisites

Verify that you have configured images, flavors, block storage, and networks, and that they are available.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Images**.

The Images page lists the snapshots available to the current user.

- 4 In the Actions column of the snapshot, click **Launch**.
- 5 On the **Details** tab of the Launch Instance dialog box, configure the instance.

| Setting              | Description  |
|----------------------|--|
| Availability Zone    | By default, this value is set to the availability zone that the cloud provider provides, for example, nova.  |
| Instance Name        | Assign a name to the VM. This value is a label and is not validated. When you create an instance, a UUID is assigned to the instance. When you view the VM in vSphere, you can identify it by the UUID but not by the instance name. |
| Flavor               | Specify the size of the instance to start. The cloud administrator defines and manages flavors .   |
| Instance Count       | To start multiple instances, enter a value greater than 1. The default is 1.   |
| Instance Boot Source | Select <b>Boot from snapshot</b> , and select the snapshot from the list.  |

- 6 On the **Access & Security** tab of the Launch Instance dialog box, configure access and security parameters by specifying a key pair and security group.

| Setting         | Description  |
|-----------------|--|
| Key Pair        | Specify a key pair.<br>If the image uses a static root password or a static key set, you do not need to provide a key pair to launch the instance. A best practice is to use a key pair.   |
| Security Groups | Select the security groups to assign to the instance.<br>Security groups are sets of rules that determine which incoming network traffic is forwarded to instances. If you did not create security groups, you can assign only the default security group to the instance. |

- 7 On the **Networking** tab of the Launch Instance dialog box, click the **+** icon in the Available Networks field to add a network to the instance.
- 8 (Optional) On the **Post-Creation** tab, specify a customization script that runs after the instance starts.

- 9 In the **Advanced Options** tab, select the type of disk partition from the drop-down menu.

| Setting   | Description  |
|-----------|--|
| Automatic | The entire disk is a single partition and automatically resizes. |
| Manual    | Enables faster build times but requires manual partitioning.     |

- 10 Click **Launch**.

The new instance starts on a node in the Compute cluster.

- 11 To view the new instance, select **Project > Compute > Instances**.

The **Instances** tab shows the instance name, its private and public IP addresses, size, status, task, and power state.

## Connect to an Instance by Using SSH

To use SSH to connect to your instance, use the downloaded keypair file.

### Procedure

- 1 Copy the IP address for your instance.
- 2 Use the `ssh` command to make a secure connection to the instance.

For example:

```
$ ssh -i MyKey.pem demo@10.0.0.2
```

- 3 At the prompt, enter **yes**.

## Track Instance Use

You can track use for instances in each project. You can view instance metrics such as number of vCPUs, disks, RAM, and uptime.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Overview**.

The Overview page shows use and limit information. You can also limit the information to a specific period of time lists and download a summary in the CSV format.

## Create a Snapshot from an Instance

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

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



**Procedure**

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Instances**.

The Instances page lists the instances available to the current user.

- 4 In the Actions column, click **Create Snapshot**.

The snapshot appears on the Images page.

## 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 the *VMware Integrated OpenStack Administrator Guide*.

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

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

**Prerequisites**

- Verify with the cloud administrator that the intended filter configuration does not conflict with any 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](http://docs.openstack.org/user-guide/common/cli_install_openstack_command_line_clients.html).

**Procedure**

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

- a Create a server group with the affinity policy:

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

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

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

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

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

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

- 6 (Optional) Ask your cloud administrator to 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 with the cloud administrator that your anti-affinity filter configuration does not conflict with any 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](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 <b>affinity</b> or <b>anti-affinity</b> . |

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

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

- 3 (Optional) Ask your cloud administrator to confirm that the new rule and the server group instances are created 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.

# Working with Volumes

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

You can attach a volume to a running instance or detach a volume and attach it to another instance at any time. You can also create a snapshot from or delete a volume.

Only administrative users can create volume types.

This chapter includes the following topics:

- [Create a Volume](#)
- [Modify Existing Volumes](#)
- [Delete Existing Volumes](#)
- [Attach a Volume to an Instance](#)
- [Detach a Volume](#)
- [Create a Snapshot from a Volume](#)

## Create a Volume

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

### Prerequisites

Upload an image for the volume. See [Chapter 2 Managing Images for the Image Service](#).

### Procedure

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

The Volume & Snapshots page lists the volumes currently configured that are available to the current user.

- 4 Click **Create Volume**.

## 5 Create the volume.

| Option      | Description  |
|-------------|--|
| Volume Name | Enter a name for the new volume.                   |
| Description | (Optional) Enter a description for the new volume. |
| Type        | Leave blank.                                       |
| Size        | Enter the size of the volume.                      |

## 6 Specify the volume source.

| Option                  | Description  |
|-------------------------|--|
| No source, empty volume | Creates an empty volume. An empty volume does not contain a file system or a partition table.  |
| Snapshot                | Creates a volume from a snapshot. If you choose this option, the <b>Use snapshot as a source</b> field appears. Select the snapshot from the list. The options to use a snapshot or a volume as the source for a volume appear only if snapshots or volumes exist.             |
| Image                   | Select this option to create a volume from an image. If you choose this option, the <b>Use image as a source</b> field appears. Select the image from the list.  |
| Availability Zone       | Select the Availability Zone from the list. By default, this value is set to the availability zone specified by the cloud provider, for example, <b>us-west</b> or <b>apac-south</b> . The default can also be <b>nova</b> .   |
| Volume                  | Creates a volume from an existing volume. If you choose this option, the <b>Use volume as a source field</b> appears. You can select the volume from the list. The options to use a snapshot or a volume as the source for a volume appear only if snapshots or volumes exist. |

## 7 Click **Create Volume** at the bottom of the page.

The Volume & Snapshots page appears again, showing the new volume in the table.

## Modify Existing Volumes

You can modify the name and description for an existing volume. When you delete an instance, the attached volumes and their data are not destroyed.

### Procedure

- 1 Go to the Volumes page and locate the volume to modify.
- 2 In the Actions column, click **Edit Volume**.
- 3 Modify the settings and click **Edit Volume**.

## Delete Existing Volumes

When you delete an instance, the attached volumes and their data are not destroyed

### Procedure

- 1 Go to the Volumes page and select the volume to delete.
- 2 Select the volumes to delete.
- 3 Click **Delete Volumes**.
- 4 When prompted, confirm the deletion.

The deleted volume no longer appears on the Volumes page.

## Attach a Volume to an Instance

After you create one or more volumes, you can attach them to instances. You can attach a volume to one instance at a time.

### Procedure

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

The Volume & Snapshots page lists the volumes currently available to the current user.

- 4 Select the volume to add to an instance and select **More > Edit Attachments** in the Actions column.
- 5 From the **Attach to Instance** drop-down menu, select the instance to which you want to attach the volume.
- 6 Click **Attach Volume**.

The new volume appears in the list of available volumes.

## Detach a Volume

You can detach a volume from one instance and attach it to another.

### Procedure

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

The Volume & Snapshots page lists the volumes currently available to the current user.

- 4 Select the volume to detach and click **Edit Attachments**.
- 5 Click **Detach Volume**.
- 6 Confirm the action at the prompt.

The volume is now available and can be attached to a different instance.

## Create a Snapshot from a Volume

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

### Prerequisites

Detach the volume from the instance before you take the snapshot. Creating a snapshot from an attached volume is not allowed.

### Procedure

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

The Volume & Snapshots page lists the volumes currently configured that are available to the current user.

- 4 Select the volume to add to an instance and select **More > Create Snapshot** in the Actions column.
- 5 Enter a snapshot name and optional description.
- 6 Click **Create Volume Snapshot**.

The Volume & Snapshots page appears again, showing the new snapshot in the table on the **Volume Snapshots** tab.

# Working with Orchestration and Stacks



You can use the OpenStack Orchestration service to orchestrate multiple composite cloud applications. It supports the native OpenStack Heat Orchestration Template (HOT) format through a REST API, and the Amazon Web Services (AWS) CloudFormation template format through a Query API that is compatible with CloudFormation.

You use templates to create stacks. A stack configures the automated creation of most OpenStack resource types, including instances, floating IP addresses, volumes, security groups, and users.

With orchestration templates, application developers can define the parameters for automating the deployment of infrastructure, services, and applications. Templates are static files that you can use directly for creating a stack.

You can also create a stack that combines a template with an environment file. An environment file supplies a unique set of values to the parameters defined by the template. By using environment files with templates, you can create many unique stacks from a single template.

For information about how to create template and environment files, see the OpenStack documentation at [http://docs.openstack.org/developer/heat/template\\_guide/index.html](http://docs.openstack.org/developer/heat/template_guide/index.html).

This chapter includes the following topics:

- [Start a New Orchestration Stack](#)
- [Modify an Orchestration Stack](#)
- [Delete an Orchestration Stack](#)

## Start a New Orchestration Stack

With orchestration stacks, you can launch and manage multiple composite cloud applications. You start a new stack by specifying the template and environment files, and defining other operational settings, including user credentials, database access settings, and the Linux distribution.

### Prerequisites

Verify that the template and environment file for the stack are created and available. For information about creating template and environment files, see the OpenStack documentation at [http://docs.openstack.org/developer/heat/template\\_guide/index.html](http://docs.openstack.org/developer/heat/template_guide/index.html).



**Procedure**

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Orchestration > Stacks**.

The Stacks page lists the stacks available to the current user.

- 4 Click **Launch Stack**.
- 5 Select the template for the new stack.

| Option                          | Description  |
|---------------------------------|--|
| Template Source                 | Select the template source: URL, File, or Direct Input.  |
| Template URL or File or Data    | Dynamically changes depending on what you select for Template Source. Enter the URL, browse to the file location, or paste the template text.    |
| Environment Source              | Select the environment source: URL, File, or Direct Input.   |
| Environment URL or File or Data | Dynamically changes depending on what you select for Environment Source. Enter the URL, browse to the file location, or paste the template text. |

- 6 Click **Next**.
- 7 Configure the new stack.

| Option                     | Description  |
|----------------------------|--|
| Stack Name                 | Name to identify the stack.  |
| Creation Timeout (minutes) | Number of minutes before the launch of the stack times out.              |
| Rollback On Failure        | Select this check box to roll back changes if the stack fails to launch. |
| Password for user "demo"   | Password for the default user after the stack is created.                |
| DBUsername                 | Name of the database user.   |
| Linux Distribution         | Linux distribution that is used in the stack.                            |
| DB Root Password           | Root password for the database.  |
| Key Name                   | Key pair for logging into the stack.                                     |
| DB Name                    | Name of the database.  |
| DB Password                | Password for the database.   |
| Instance Type              | Flavor for the instance.   |

- 8 Click **Launch** to create the stack.
- 9 (Optional) Verify that the new stack appears on the Stacks page.
- 10 (Optional) Click the stack to view the stack details.

| Detail   | Description                          |
|----------|--------------------------------------|
| Topology | Visual topology of the stack.        |
| Overview | Parameters and details of the stack. |

| Detail    | Description                    |
|-----------|--------------------------------|
| Resources | Resources that the stack uses. |
| Events    | Events related to the stack.   |

## Modify an Orchestration Stack

You can modify a stack by updating the template file, environment file, or stack parameters.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Orchestration > Stacks**.  
The Stacks page lists the stacks available to the current user.
- 4 Select the stack to update.
- 5 Click **Change Stack Template**.
- 6 (Optional) In the Select Template dialog box, modify the template or environment file selection.
- 7 Click **Next**.
- 8 (Optional) In the Update Stack Parameters dialog box, modify the parameter values.
- 9 Click **Update**.
- 10 (Optional) On the Stacks page, verify that the changes to the stack configuration are applied.

## Delete an Orchestration Stack

When you delete a stack, you also delete the resources that that stack generates.

### Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select the project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Orchestration > Stacks**.  
The Stacks page lists the stacks available to the current user.
- 4 Select the stack to delete and click **Delete Stack**.
- 5 Confirm the action at the prompt.
- 6 (Optional) Verify that the deleted stack no longer appears on the Stacks page.