

VMware Integrated OpenStack User's Guide

Update 2

Modified on 13 NOV 2018

VMware Integrated OpenStack 4.1



vmware®

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

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

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

docfeedback@vmware.com

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

Copyright © 2015–2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

- 1 VMware Integrated OpenStack User's Guide 4**
- 2 Updated Information 5**
- 3 Log In to the VMware Integrated OpenStack Dashboard 6**
- 4 Working with Images 8**
 - [Import Images Using the GUI 8](#)
 - [Import Images Using the CLI 9](#)
 - [Import Images in Unsupported Formats 10](#)
- 5 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](#)
- 6 Working with Networks 18**
 - [Create a Network 18](#)
 - [Create a Router 19](#)
 - [Configuring LBaaS v2.0 20](#)
- 7 Working with Instances in OpenStack 25**
 - [Start an OpenStack Instance from an Image 25](#)
 - [Start an OpenStack Instance from a Snapshot 27](#)
 - [Connect to an Instance by Using SSH 28](#)
 - [Track Instance Use 28](#)
 - [Create a Snapshot from an Instance 28](#)
 - [Using Affinity and Anti-Affinity to Place OpenStack Instances 29](#)
- 8 Working with Volumes 32**
 - [Create a Volume 32](#)
 - [Transfer a Volume 33](#)
- 9 Working with Orchestration and Stacks 35**
 - [Start a New Orchestration Stack 35](#)
 - [Modify an Orchestration Stack 37](#)
 - [Delete an Orchestration Stack 37](#)

VMware Integrated OpenStack User's Guide

1

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

Intended Audience

This guide is for cloud users who want to work 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 User's Guide* is updated with each release of the product or when necessary.

This table provides the update history of the *VMware Integrated OpenStack User's Guide*.

Revision	Description
Update 2 (13 NOV 2018)	<ul style="list-style-type: none">■ Added LBaaS section.■ Various corrections and improvements.
Update 1 (08 OCT 2018)	Various corrections and improvements.
18 JAN 2018	Initial release.

Log In to the VMware Integrated OpenStack Dashboard

3

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 3-1. VMware Integrated OpenStack Overview Page

vmware® service writer_andy Sign Out

Project Compute Overview

Overview

Limit Summary

- Instances Used 0 of 10
- VCPUs Used 0 of 20
- RAM Used 0Bytes of 50.0GB
- Floating IPs Used 0 of 50
- Security Groups Used 1 of 10
- Volumes Used 0 of 10
- Volume Storage Used 0Bytes of 1000.0GB

Usage Summary

Select a period of time to query its usage:

From: 2014-12-01 To: 2014-12-22 [Submit](#) The date should be in YYYY-mm-dd format.

Active Instances: 0 Active RAM: 0Bytes This Period's VCPU-Hours: 0.00 This Period's GB-Hours: 0.00

Usage

[Download CSV Summary](#)

Instance Name	VCPUs	Disk	RAM	Uptime
No items to display.				
Displaying 0 items				

Waiting for 10.146.30.250...

Working with Images

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

If you have images in vSphere that you want to use in OpenStack, you can export them in one of the supported formats and upload them to the image service. You can also use the `glance-import` tool to convert RAW, QCOW2, VDI, and VHD images to the VMDK format and use them in OpenStack.

This chapter includes the following topics:

- [Import Images Using the GUI](#)
- [Import Images Using the CLI](#)
- [Import Images in Unsupported Formats](#)

Import Images Using the GUI

You can import images in the VMware Integrated OpenStack dashboard.

The following image formats are supported:

- VMDK
- ISO
- OVA

To upload images in another format, see [Import Images in Unsupported Formats](#).

Note ISO images cannot be used to create volumes.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select your project from the drop-down menu in the title bar.
- 3 Select **Project > 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.
Source Type	Select File to select a local file or URL to specify a remote file.
Format	Select ISO or VMDK . For images in OVA format, select VMDK as the disk format.
Disk Adapter Type	For VMDK images, select the adapter type.
Minimum Disk (GB)	Specify the minimum disk size for the image in gigabytes.
Minimum RAM (MB)	Specify the minimum RAM for the image in megabytes.
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

You can launch OpenStack instances using the imported image. See [Start an OpenStack Instance from an Image](#).

In the **Actions** column next to an image, you can also 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

To upload images in another format, see [Import Images in Unsupported Formats](#).

Note ISO images cannot be used to create volumes.

Procedure

- 1 Log in to the OpenStack Management Server as `viouser`.
- 2 Load the credentials file for your user account.

```
source user-credentials.rc
```

3 Create the image in Glance.

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

Option	Description
<code>--name</code>	Enter a name for the image file in the image service.
<code>--file</code>	Enter the path to the desired image file.
<code>--disk_format</code>	Enter the disk format of the source image. You can specify <code>iso</code> or <code>vmdk</code> . For images in OVA format, use <code>vmdk</code> as the disk format.
<code>--container_format</code>	Enter <code>bare</code> . The container format argument is not currently used by Glance.
<code>--property vmware_adaptype</code>	Specify the adapter type of the VMDK disk. If you do not include this parameter, the adapter type is determined by introspection. Note <ul style="list-style-type: none"> ■ For disks using paravirtual adapters, include this parameter and set it to <code>paraVirtual</code>. ■ For disks using LSI Logic SAS adapters, include this parameter and set it to <code>lsiLogicsas</code>.
<code>--property vmware_disktype</code>	Specify <code>sparse</code> , <code>preallocated</code> , or <code>streamOptimized</code> . If you do not include this parameter, the disk type is determined by introspection.
<code>--property vmware_ostype</code>	Specify the operating system on the image.

What to do next

You can launch OpenStack instances using the imported image. See [Start an OpenStack Instance from an Image](#). You can also run the `openstack image list` command to see all images in your project.

Import Images in Unsupported Formats

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

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

Procedure

- 1 Log in to the OpenStack Management Server as `viouser`.
- 2 Log in to the controller node as `viouser`.
- 3 Load the credentials file for your user account.

```
source user-credentials.rc
```

4 Import the image.

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

Parameter	Description
--name	Enter a name for the image file in the image service.
--url	Enter the URL where the source image is located.
--image-format	Specify the format of the source image file. Non-VMDK images are converted automatically to the VMDK format. You can use the following formats: <ul style="list-style-type: none"> ▪ VMDK ▪ OVA ▪ RAW ▪ QCOW2 ▪ VDI ▪ VHD

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

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

What to do next

You can launch OpenStack instances using the imported image. See [Start an OpenStack Instance from an Image](#). You can also run the `glance image-list` command to see all images in your project.

Configuring Access and Security for Instances

5

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

- | | |
|------------------------|---|
| Security groups | 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. |
| Key pairs | 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. |
| Floating IPs | 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 NSX Transformers, 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.

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

- 8 Click **Add**.

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

Enabling SSH and ICMP Access

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

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 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**.
- 6 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.

- 7 Click **Associate**.
- 8 (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. .
- 9 To release the floating IP address back into the pool of addresses, click **More** and select **Release Floating IP**.
- 10 Click the **Floating IPs** tab and select the IP address.
- 11 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](#)
- [Configuring LBaaS v2.0](#)

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
Create Subnet	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 Create Subnet .
Subnet Name	(Optional) Enter a name for the subnet.
Network Address	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.
IP Version	Select IPv4 or IPv6 from the drop-down menu.
Gateway IP	Enter the IP address for a specific gateway.
Disable Gateway	(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
Enable DHCP	(Optional) Select this option to enable DHCP. Consult with your network administrator.
Allocation Pools	Specify IP address pools for use by devices in the new network.
DNS Name Servers	Specify DNS servers for the new network.
Host Routes	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.

Configuring LBaaS v2.0

You can enable LBaaS v2.0 to distribute incoming requests among designated instances.

Load balancer as a service (LBaaS) v2.0 gives you the ability to create load balancers on demand, ensuring that workloads are shared predictably among instances and system resources are used more effectively.

The LBaaS configuration process also creates a health monitor and associates it with the LBaaS pool. The health monitor is a Neutron service that checks whether the instances are still running on the specified protocol and port.

You can enable LBaaS v2.0 on VMware Integrated OpenStack deployments with NSX-T or NSX-V networking.

Note The `admin_state` parameter for LBaaS pools is not supported on NSX-V deployments, and setting the admin state of a pool to down has no effect. To prevent network traffic from reaching the members of a pool, set the admin state of each member to down.

Configure LBaaS Using the CLI

You can configure LBaaS using the command line interface on the node running OpenStack Neutron.

Prerequisites

- Create a public subnet and router on your network. For an NSX-V deployment, the router type must be `exclusive`.

Note You can create the load balancer on a tenant subnet, but you must assign it a floating IP address.

- Configure at least one client and at least two server instances.

Procedure

- 1 Log in to the node running OpenStack Neutron and switch to the root user.

```
sudo su -
```

- 2 Create a load balancer.

```
neutron lbaas-loadbalancer-create --name lb-name lb-subnet-id
```

Only members of the specified subnet can be added to the LBaaS pool.

- 3 Create a listener for the new load balancer.

```
neutron lbaas-listener-create --loadbalancer lb-name --protocol {HTTP | TCP} --protocol-port port-num --name listener-name
```

- 4 Create an LBaaS pool.

```
neutron lbaas-pool-create --lb-algorithm lb-method --listener listener-name --protocol {TCP | HTTP} --name pool-name
```

The `--lb-algorithm` parameter accepts the following values.

Argument	Description
LEAST_CONNECTIONS	New client requests are sent to the server with the fewest connections.
ROUND_ROBIN	Each server is used in turn according to the weight assigned to it.
SOURCE_IP	All connections that originate from the same source IP address are handled by the same member of the pool.

- 5 Add at least two server instances to the LBaaS pool that you created.

```
neutron lbaas-member-create --subnet lb-subnet-id --address server1-ip --protocol-port 80 pool-name
neutron lbaas-member-create --subnet lb-subnet-id --address server2-ip --protocol-port 80 pool-name
```

- 6 Set up the health monitor.

```
neutron lbaas-healthmonitor-create --delay delay-seconds --type {HTTP | TCP | PING} --max-retries number --timeout timeout-seconds --pool pool-name
```

Parameter	Description
<code>--delay</code>	Enter the time in seconds between sending probes to members.
<code>--type</code>	Specify HTTP, TCP, or PING.
<code>--max-retries</code>	Enter the number of connection failures allowed before changing the member status to INACTIVE.

Parameter	Description
<code>--timeout</code>	Enter the time in seconds that a monitor will wait for a connection to be established before it times out. The timeout value must be less than the delay value.
<code>--pool</code>	Specify the LBaaS pool that you created.

- 7 If you created the load balancer on a tenant subnet, associate a floating IP address with the load balancer.
- 8 (Optional) Send test requests to validate your LBaaS configuration.
 - a Log in to the OpenStack Management Server.
 - b Create a test `index.html` file.
 - c In the same directory, start a web server.

```
sudo python -m SimpleHTTPServer 80
```

- d Log in to the client instance.
- e Run the `wget` command and view whether your requests are being correctly load-balanced across the servers in the pool.

```
wget -O - http://mgmt-server-ip
```

Configure LBaaS Using the GUI

You can configure LBaaS using the **Create a Load Balancer** wizard in the VMware Integrated OpenStack user interface.

Prerequisites

- Create a public subnet and router on your network. For an NSX-V deployment, the router type must be exclusive.

Note You can create the load balancer on a tenant subnet, but you must assign it a floating IP address.

- Configure at least one client and at least two server instances.

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard and select **Project > Network > Load Balancers**.
- 2 Click **Create Load Balancer**.
- 3 On the **Load Balancer Details** page, specify a subnet.
Only members of this subnet can be added to the LBaaS pool.

- 4 Click **Next**. On the **Listener Details** page, create a listener for the new load balancer.
- 5 Click **Next**. On the **Pool Details** page, create an LBaaS pool.

Supported load balancing methods are described as follows:

Method	Description
LEAST_CONNECTIONS	New client requests are sent to the server with the fewest connections.
ROUND_ROBIN	Each server is used in turn according to the weight assigned to it.
SOURCE_IP	All connections that originate from the same source IP address are handled by the same member of the pool.

- 6 Click **Next**. On the **Pool Members** page, select the server and client instances to add to the load balancer pool.
- 7 Click **Next**. On the **Monitor Details** page, set up the health monitor.

Parameter	Description
Monitor type	Specify HTTP , PING , or TCP .
Interval	Enter the time in seconds between sending probes to members.
Retries	Enter the number of connection failures allowed before changing the member status to INACTIVE .
Timeout	Enter the time in seconds that a monitor will wait for a connection to be established before it times out. The timeout value must be less than the interval value.

If you select **HTTP**, you must also configure the HTTP method, expected status code, and URL.

- 8 If you created the load balancer on a tenant subnet, associate a floating IP address with the load balancer.
 - a Click the down arrow to the right of the load balancer and select **Associate Floating IP**.
 - b Select a floating IP address or pool and click **Associate**.
- 9 (Optional) Send test requests to validate your LBaaS configuration.
 - a Log in to the OpenStack Management Server.
 - b Create a test `index.html` file.
 - c In the same directory, start a web server.

```
sudo python -m SimpleHTTPServer 80
```

- d Log in to the client instance.
- e Run the `wget` command to view whether your requests are being correctly load-balanced across the servers in the pool.

```
wget -O - http://mgmt-server-ip
```

What to do next

You can open the load balancer and click **Create Listener** to add listeners to it.

Working with Instances in OpenStack

7

Instances are virtual machines that run in the cloud. You can launch instances, track their usage, and create snapshots.

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: nova .
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 1 .
Instance Boot Source	Select Boot from image , 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. 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. 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 Boot from snapshot , 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. 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. 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 [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.

Working with Volumes

Volumes are block storage devices that provide persistent storage for instances.

After you create a volume, you can attach it to a running instance. You can later detach the volume and attach it to a different instance. You can also create a snapshot of a volume, launch an instance from it, and upload it to Glance as an image.

This chapter includes the following topics:

- [Create a Volume](#)
- [Transfer a Volume](#)

Create a Volume

You create volumes and attach them to instances to provide persistent storage.

Prerequisites

If you want to create a volume from an image, upload the desired image. See [Chapter 4 Working with Images](#).

Procedure

- 1 Log in to the VMware Integrated OpenStack dashboard.
- 2 Select your project from the drop-down menu in the title bar.
- 3 Select **Project > Compute > Volumes** and click **Create Volume**.
- 4 Configure the volume.

Option	Description
Volume Name	Enter a name for the new volume.
Description	Enter a description for the volume.
Volume Source	Select No source, empty volume, Snapshot, Image, or Volume . If you select Snapshot, Image, or Volume , specify the desired object from the next drop-down list.
Type	If you selected No source, empty volume or Image as the volume source, select a volume type for the volume. For volumes whose source is a volume snapshot or another volume, the volume type is inherited from the source.

Option	Description
Size (GiB)	Enter the size of the volume in gibibytes.
Availability Zone	If you selected No source, empty volume or Image as the volume source, specify the availability zone in which to create the volume. For volumes whose source is a volume snapshot or another volume, the availability zone is inherited from the source.

5 Click **Create Volume**.

What to do next

In the **Actions** column to the right of the volume, you can perform the following actions:

- Click **Edit Volume** to modify the name and description of the volume and whether it is bootable.
- Click **Extend Volume** to increase the size of an unattached volume.
- Click **Launch as Instance** to create an instance using an unattached volume.
- Click **Manage Attachments** to attach the volume to or detach the volume from an instance.
- Click **Create Snapshot** to take a snapshot of the volume.

Note Creating a snapshot of a volume attached to an instance can result in a corrupted snapshot. If possible, detach the volume before creating the snapshot.

- Click **Change Volume Type** to modify the volume type and migration policy.
- Click **Upload to Image** to upload the volume to Glance as an image.
- Click **Create Transfer** to assign ownership of an unattached volume to a different project. For details, see [Transfer a Volume](#).
- Click **Delete Volume** to delete an unattached volume.
- Click **Update Metadata** to add, remove, or change volume metadata.

Transfer a Volume

You can assign ownership of an unattached volume to another project.

Prerequisites

Ensure that the volume that you want to transfer is not attached to an instance.

Procedure

- To initiate a transfer, perform the following steps:
 - a Log in to the VMware Integrated OpenStack dashboard.
 - b Select your project from the drop-down menu in the title bar.
 - c Select **Project > Compute > Volumes**.
 - d In the **Actions** column next to the volume that you want to transfer, click **Create Transfer**.

- e Enter a name for the transfer task and click **Create Volume Transfer**.
- f Record or download the transfer ID and authorization key displayed on the **Volume Transfer Details** page and send this information to the user who will accept the transfer.

Important After you close the **Volume Transfer Details** page, the transfer ID and authorization key can no longer be retrieved. If the transfer ID or authorization key are lost, you must cancel the transfer and initiate it again.

- To receive a transfer, perform the following steps:
 - a Log in to the VMware Integrated OpenStack dashboard.
 - b Select your project from the drop-down menu in the title bar.
 - c Select **Project > Compute > Volumes** and click **Accept Transfer**.
 - d Enter the transfer ID and authorization key that you received from the user who initiated the transfer.
 - e Click **Accept Volume Transfer**.

Working with Orchestration and Stacks

9

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.

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.

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.