Red Hat OpenStack Platform 14

Instances and Images Guide

Managing Instances and Images

Red Hat OpenStack Platform 14 Instances and Images Guide

Managing Instances and Images

OpenStack Team
rhos-docs@redhat.com

Legal Notice

Abstract

The Instances and Images guide provides procedures for the management of instances, images of a Red Hat OpenStack Platform environment.

Table of Contents

PREFACE

Red Hat OpenStack Platform (Red Hat OpenStack Platform) provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It offers a massively scalable, fault-tolerant platform for the development of cloud-enabled workloads.

This guide discusses procedures for creating and managing images, and instances. It also mentions the procedure for configuring the storage for instances for Red Hat OpenStack Platform.

You can manage the cloud using either the OpenStack dashboard or the command-line clients. Most procedures can be carried out using either method; some of the more advanced procedures can only be executed on the command line. This guide provides procedures for the dashboard where possible.

> **NOTE**
>
> For the complete suite of documentation for Red Hat OpenStack Platform, see Red Hat OpenStack Platform Documentation Suite.

CHAPTER 1. IMAGE SERVICE

This chapter discusses the steps you can follow to manage images and storage in Red Hat OpenStack Platform.

A virtual machine image is a file which contains a virtual disk which has a bootable operating system installed on it. Virtual machine images are supported in different formats. The following formats are available on Red Hat OpenStack Platform:

- **RAW** – Unstructured disk image format.

- **QCOW2** – Disk format supported by QEMU emulator. This format includes QCOW2v3 (sometimes referred to as QCOW3), which requires QEMU 1.1 or higher.

- **ISO** – Sector-by-sector copy of the data on a disk, stored in a binary file.

- **AKI** – Indicates an Amazon Kernel Image.

- **AMI** – Indicates an Amazon Machine Image.

- **ARI** – Indicates an Amazon RAMDisk Image.

- **VDI** – Disk format supported by VirtualBox virtual machine monitor and the QEMU emulator.

- **VHD** – Common disk format used by virtual machine monitors from VMware, VirtualBox, and others.

- **VMDK** – Disk format supported by many common virtual machine monitors.

While **ISO** is not normally considered a virtual machine image format, since ISOs contain bootable filesystems with an installed operating system, you can treat them the same as you treat other virtual machine image files.

To download the official Red Hat Enterprise Linux cloud images, your account must have a valid Red Hat Enterprise Linux subscription:

- Red Hat Enterprise Linux 7 KVM Guest Image

- Red Hat Enterprise Linux 6 KVM Guest Image

You will be prompted to enter your Red Hat account credentials if you are not logged in to the Customer Portal.

1.1. UNDERSTANDING THE IMAGE SERVICE

The following notable OpenStack Image service (glance) features are available.

1.1.1. Image Signing and Verification

Image signing and verification protects image integrity and authenticity by enabling deployers to sign images and save the signatures and public key certificates as image properties.

By taking advantage of this feature, you can:

- Sign an image using your private key and upload the image, the signature, and a reference to your public key certificate (the verification metadata). The Image service then verifies that the signature is valid.

- Create an image in the Compute service, have the Compute service sign the image, and upload the image and its verification metadata. The Image service again verifies that the signature is valid.

- Request a signed image in the Compute service. The Image service provides the image and its verification metadata, allowing the Compute service to validate the image before booting it.

For information on image signing and verification, refer to the Validate Glance Images chapter of the *Manage Secrets with OpenStack Key Manager Guide* .

1.1.2. Image Conversion

Image conversion converts images by calling the task API while importing an image.

As part of the import workflow, a plugin provides the image conversion. This plugin can be activated or deactivated based on the deployer configuration. Therefore, the deployer needs to specify the preferred format of images for the deployment.

Internally, the Image service receives the bits of the image in a particular format. These bits are stored in a temporary location. The plugin is then triggered to convert the image to the target format, and moved to a final destination. When the task is finished, the temporary location is deleted. As a result, the format uploaded initially is not retained by the Image service.

> **NOTE**
>
> The conversion can be triggered only when **importing** an image (the old copy-from). It does not run when **uploading** an image. For example:
>
> ```
> $ glance task-create --type import --input '{"import_from_format": "qcow2",
> "import_from": "http://127.0.0.1:8000/test.qcow2", "image_properties": {"disk_format":
> "qcow2", "container_format": "bare"}}'
> ```

1.1.2.1. Enabling Image Conversion

To enable image conversion, create an environment file that contains the following parameter value and include the new environment file with the **-e** option in the **openstack overcloud deploy** command:

```
parameter_defaults:
  GlanceImageImportPlugins:'image_conversion'
```

1.1.3. Image Introspection

Every image format comes with a set of metadata embedded inside the image itself. For example, a stream optimized **vmdk** would contain the following parameters:

```
$ head -20 so-disk.vmdk

# Disk DescriptorFile
version=1
CID=d5a0bce5
```

```
parentCID=ffffffff
createType="streamOptimized"

# Extent description
RDONLY 209714 SPARSE "generated-stream.vmdk"

# The Disk Data Base
#DDB

ddb.adapterType = "buslogic"
ddb.geometry.cylinders = "102"
ddb.geometry.heads = "64"
ddb.geometry.sectors = "32"
ddb.virtualHWVersion = "4"
```

By introspecting this *vmdk*, you can easily know that the *disk_type* is *streamOptimized*, and the *adapter_type* is *buslogic*. These metadata parameters are useful for the consumer of the image. In Compute, the workflow to instantiate a *streamOptimized* disk is different from the one to instantiate a *flat* disk. This new feature allows metadata extraction. You can achieve image introspection by calling the task API while importing the image. An administrator can override metadata settings.

1.1.4. Interoperable Image Import

The OpenStack Image service provides two methods for importing images using the interoperable image import workflow:

- **web-download** (default) for importing images from a URI and

- **glance-direct** for importing from a local file system.

1.2. MANAGE IMAGES

The OpenStack Image service (glance) provides discovery, registration, and delivery services for disk and server images. It provides the ability to copy or snapshot a server image, and immediately store it away. Stored images can be used as a template to get new servers up and running quickly and more consistently than installing a server operating system and individually configuring services.

1.2.1. Create an Image

This section provides you with the steps to manually create OpenStack-compatible images in the QCOW2 format using Red Hat Enterprise Linux 7 ISO files, Red Hat Enterprise Linux 6 ISO files, or Windows ISO files.

1.2.1.1. Use a KVM Guest Image With Red Hat OpenStack Platform

You can use a ready RHEL KVM guest QCOW2 image:

- RHEL 7.2 KVM Guest Image

- RHEL 6.8 KVM Guest Image

These images are configured with **cloud-init** and must take advantage of ec2-compatible metadata services for provisioning SSH keys in order to function properly.

Ready Windows KVM guest QCOW2 images are not available.

NOTE

For the KVM guest images:

- The **root** account in the image is disabled, but **sudo** access is granted to a special user named **cloud-user**.

- There is no **root** password set for this image.

The **root** password is locked in **/etc/shadow** by placing **!!** in the second field.

For an OpenStack instance, it is recommended that you generate an ssh keypair from the OpenStack dashboard or command line and use that key combination to perform an SSH public authentication to the instance as root.

When the instance is launched, this public key will be injected to it. You can then authenticate using the private key downloaded while creating the keypair.

If you do not want to use keypairs, you can use the **admin** password that has been set using the Inject an **admin** Password Into an Instance procedure.

If you want to create custom Red Hat Enterprise Linux or Windows images, see Create a Red Hat Enterprise Linux 7 Image, Create a Red Hat Enterprise Linux 6 Image , or Create a Windows Image .

1.2.1.2. Create Custom Red Hat Enterprise Linux or Windows Images

Prerequisites:

- Linux host machine to create an image. This can be any machine on which you can install and run the Linux packages.

- libvirt, virt-manager (run command **yum groupinstall -y @virtualization**). This installs all packages necessary for creating a guest operating system.

- Libguestfs tools (run command **yum install -y libguestfs-tools-c**). This installs a set of tools for accessing and modifying virtual machine images.

- A Red Hat Enterprise Linux 7 or 6 ISO file (see RHEL 7.2 Binary DVD or RHEL 6.8 Binary DVD) or a Windows ISO file. If you do not have a Windows ISO file, visit the Microsoft TechNet Evaluation Center and download an evaluation image.

- Text editor, if you want to change the **kickstart** files (RHEL only).

NOTE

In the following procedures, all commands with the **[root@host]#** prompt should be run on your host machine.

1.2.1.2.1. Create a Red Hat Enterprise Linux 7 Image

This section provides you with the steps to manually create an OpenStack-compatible image in the QCOW2 format using a Red Hat Enterprise Linux 7 ISO file.

1. Start the installation using **virt-install** as shown below:

   ```
   [root@host]# qemu-img create -f qcow2 rhel7.qcow2 8G
   ```

```
[root@host]# virt-install --virt-type kvm --name rhel7 --ram 2048 \
--cdrom /tmp/rhel-server-7.2-x86_64-dvd.iso \
--disk rhel7.qcow2,format=qcow2 \
--network=bridge:virbr0 --graphics vnc,listen=0.0.0.0 \
--noautoconsole --os-type=linux --os-variant=rhel7
```

This launches an instance and starts the installation process.

> **NOTE**
>
> If the instance does not launch automatically, run the **virt-viewer** command to view the console:
>
> ```
> [root@host]# virt-viewer rhel7
> ```

2. Set up the virtual machine as follows:

   a. At the initial Installer boot menu, choose the **Install Red Hat Enterprise Linux 7.**X option.
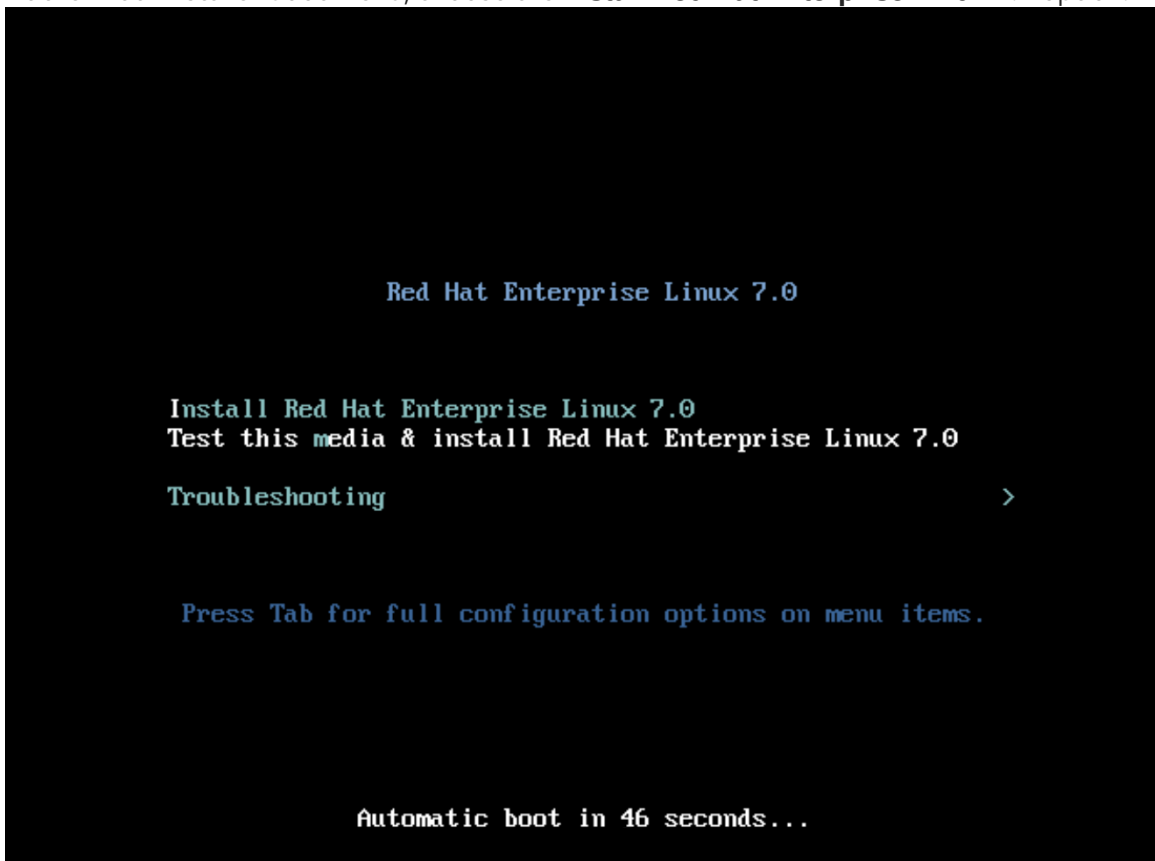
   b. Choose the appropriate **Language** and **Keyboard** options.

c. When prompted about which type of devices your installation uses, choose **Auto-detected installation media**.

d. When prompted about which type of installation destination, choose **Local Standard Disks**.

For other storage options, choose **Automatically configure partitioning**.

e. For software selection, choose **Minimal Install**.

f. For network and host name, choose **eth0** for network and choose a **hostname** for your device. The default host name is **localhost.localdomain**.

g. Choose the **root** password.

The installation process completes and the **Complete!** screen appears.

3. After the installation is complete, reboot the instance and log in as the root user.

4. Update the **/etc/sysconfig/network-scripts/ifcfg-eth0** file so it only contains the following values:

```
TYPE=Ethernet
DEVICE=eth0
ONBOOT=yes
BOOTPROTO=dhcp
NM_CONTROLLED=no
```

5. Reboot the machine.

6. Register the machine with the Content Delivery Network.

```
# sudo subscription-manager register
# sudo subscription-manager attach --pool=Valid-Pool-Number-123456
# sudo subscription-manager repos --enable=rhel-7-server-rpms
```

7. Update the system:

```
# yum -y update
```

8. Install the **cloud-init** packages:

```
# yum install -y cloud-utils-growpart cloud-init
```

9. Edit the **/etc/cloud/cloud.cfg** configuration file and under **cloud_init_modules** add:

```
- resolv-conf
```

The **resolv-conf** option automatically configures the **resolv.conf** when an instance boots for the first time. This file contains information related to the instance such as **nameservers**, **domain** and other options.

10. Add the following line to **/etc/sysconfig/network** to avoid problems accessing the EC2 metadata service:

```
NOZEROCONF=yes
```

11. To ensure the console messages appear in the **Log** tab on the dashboard and the **nova console-log** output, add the following boot option to the **/etc/default/grub** file:

```
GRUB_CMDLINE_LINUX_DEFAULT="console=tty0 console=ttyS0,115200n8"
```

Run the **grub2-mkconfig** command:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

The output is as follows:

```
Generating grub configuration file ...
Found linux image: /boot/vmlinuz-3.10.0-229.7.2.el7.x86_64
Found initrd image: /boot/initramfs-3.10.0-229.7.2.el7.x86_64.img
Found linux image: /boot/vmlinuz-3.10.0-121.el7.x86_64
Found initrd image: /boot/initramfs-3.10.0-121.el7.x86_64.img
Found linux image: /boot/vmlinuz-0-rescue-b82a3044fb384a3f9aeacf883474428b
Found initrd image: /boot/initramfs-0-rescue-b82a3044fb384a3f9aeacf883474428b.img
done
```

12. Un–register the virtual machine so that the resulting image does not contain the same subscription details for every instance cloned based on it:

```
# subscription-manager repos --disable=*
# subscription-manager unregister
# yum clean all
```

13. Power off the instance:

```
# poweroff
```

14. Reset and clean the image using the **virt-sysprep** command so it can be used to create instances without issues:

```
[root@host]# virt-sysprep -d rhel7
```

15. Reduce image size using the **virt-sparsify** command. This command converts any free space within the disk image back to free space within the host:

```
[root@host]# virt-sparsify --compress /tmp/rhel7.qcow2 rhel7-cloud.qcow2
```

This creates a new **rhel7-cloud.qcow2** file in the location from where the command is run.

The **rhel7-cloud.qcow2** image file is ready to be uploaded to the Image service. For more information on uploading this image to your OpenStack deployment using the dashboard, see Upload an Image.

### 1.2.1.2.2. Create a Red Hat Enterprise Linux 6 Image

This section provides you with the steps to manually create an OpenStack-compatible image in the QCOW2 format using a Red Hat Enterprise Linux 6 ISO file.

1. Start the installation using **virt-install**:

```
[root@host]# qemu-img create -f qcow2 rhel6.qcow2 4G
[root@host]# virt-install --connect=qemu:///system --network=bridge:virbr0 \
--name=rhel6 --os-type linux --os-variant rhel6 \
--disk path=rhel6.qcow2,format=qcow2,size=10,cache=none \
--ram 4096 --vcpus=2 --check-cpu --accelerate \
--hvm --cdrom=rhel-server-6.8-x86_64-dvd.iso
```

This launches an instance and starts the installation process.

> **NOTE**
>
> If the instance does not launch automatically, run the **virt-viewer** command to view the console:
>
> ```
> [root@host]# virt-viewer rhel6
> ```

2. Set up the virtual machines as follows:

a. At the initial Installer boot menu, choose the **Install or upgrade an existing system** option.

Step through the installation prompts. Accept the defaults.

The installer checks for the disc and lets you decide whether you want to test your installation media before installation. Select **OK** to run the test or **Skip** to proceed without testing.

b. Choose the appropriate **Language** and **Keyboard** options.

c. When prompted about which type of devices your installation uses, choose **Basic Storage Devices**.

d. Choose a **hostname** for your device. The default host name is **localhost.localdomain**.

e. Set **timezone** and **root** password.

f. Based on the space on the disk, choose the type of installation.

g. Choose the **Basic Server** install, which installs an SSH server.

h. The installation process completes and **Congratulations, your Red Hat Enterprise Linux installation is complete** screen appears.

3. Reboot the instance and log in as the **root** user.

4. Update the **/etc/sysconfig/network-scripts/ifcfg-eth0** file so it only contains the following values:

```
TYPE=Ethernet
DEVICE=eth0
ONBOOT=yes
BOOTPROTO=dhcp
NM_CONTROLLED=no
```

5. Reboot the machine.

6. Register the machine with the Content Delivery Network:

```
# sudo subscription-manager register
# sudo subscription-manager attach --pool=Valid-Pool-Number-123456
# sudo subscription-manager repos --enable=rhel-6-server-rpms
```

7. Update the system:

```
# yum -y update
```

8. Install the **cloud-init** packages:

```
# yum install -y cloud-utils-growpart cloud-init
```

9. Edit the **/etc/cloud/cloud.cfg** configuration file and under **cloud_init_modules** add:

   ```
   - resolv-conf
   ```

   The **resolv-conf** option automatically configures the **resolv.conf** configuration file when an instance boots for the first time. This file contains information related to the instance such as **nameservers**, **domain**, and other options.

10. To prevent network issues, create the **/etc/udev/rules.d/75-persistent-net-generator.rules** file as follows:

    ```
    # echo "#" > /etc/udev/rules.d/75-persistent-net-generator.rules
    ```

    This prevents **/etc/udev/rules.d/70-persistent-net.rules** file from being created. If **/etc/udev/rules.d/70-persistent-net.rules** is created, networking may not function properly when booting from snapshots (the network interface is created as "eth1" rather than "eth0" and IP address is not assigned).

11. Add the following line to **/etc/sysconfig/network** to avoid problems accessing the EC2 metadata service:

    ```
    NOZEROCONF=yes
    ```

12. To ensure the console messages appear in the **Log** tab on the dashboard and the **nova console-log** output, add the following boot option to the **/etc/grub.conf**:

    ```
    console=tty0 console=ttyS0,115200n8
    ```

13. Un–register the virtual machine so that the resulting image does not contain the same subscription details for every instance cloned based on it:

    ```
    # subscription-manager repos --disable=*
    # subscription-manager unregister
    # yum clean all
    ```

14. Power off the instance:

    ```
    # poweroff
    ```

15. Reset and clean the image using the **virt-sysprep** command so it can be used to create instances without issues:

    ```
    [root@host]# virt-sysprep -d rhel6
    ```

16. Reduce image size using the **virt-sparsify** command. This command converts any free space within the disk image back to free space within the host:

    ```
    [root@host]# virt-sparsify --compress rhel6.qcow2 rhel6-cloud.qcow2
    ```

    This creates a new **rhel6-cloud.qcow2** file in the location from where the command is run.

> **NOTE**
>
> You will need to manually resize the partitions of instances based on the image in accordance with the disk space in the flavor that is applied to the instance.

The **rhel6-cloud.qcow2** image file is ready to be uploaded to the Image service. For more information on uploading this image to your OpenStack deployment using the dashboard, see Upload an Image

1.2.1.2.3. Create a Windows Image

This section provides you with the steps to manually create an OpenStack-compatible image in the QCOW2 format using a Windows ISO file.

1. Start the installation using **virt-install** as shown below:

   ```
   [root@host]# virt-install --name=name \
   --disk size=size \
   --cdrom=path \
   --os-type=windows \
   --network=bridge:virbr0 \
   --graphics spice \
   --ram=RAM
   ```

   Replace the values of the **virt-install** parameters as follows:

   - *name* — the name that the Windows guest should have.

   - *size* — disk size in GB.

   - *path* — the path to the Windows installation ISO file.

   - *RAM* — the requested amount of RAM in MB.

   > **NOTE**
   >
   > The **--os-type=windows** parameter ensures that the clock is set up correctly for the Windows guest, and enables its Hyper-V enlightenment features.

   Note that **virt-install** saves the guest image as **/var/lib/libvirt/images/**name**.qcow2** by default. If you want to keep the guest image elsewhere, change the parameter of the **--disk** option as follows:

   ```
   --disk path=filename,size=size
   ```

   Replace *filename* with the name of the file which should store the guest image (and optionally its path); for example **path=win8.qcow2,size=8** creates an 8 GB file named **win8.qcow2** in the current working directory.

   TIP

   If the guest does not launch automatically, run the **virt-viewer** command to view the console:

   ```
   [root@host]# virt-viewer name
   ```

2. Installation of Windows systems is beyond the scope of this document. For instructions on how to install Windows, see the relevant Microsoft documentation.

3. To allow the newly installed Windows system to use the virtualized hardware, you may need to install *virtio drivers* in it. To so do, first install the **virtio-win** package on the host system. This package contains the virtio ISO image, which is to be attached as a CD-ROM drive to the Windows guest. See Chapter 8. KVM Para-virtualized (virtio) Drivers in the *Virtualization Deployment and Administration Guide* for detailed instructions on how to install the **virtio-win** package, add the virtio ISO image to the guest, and install the virtio drivers.

4. To complete the setup, download and execute Cloudbase-Init on the Windows system. At the end of the installation of Cloudbase-Init, select the **Run Sysprep** and **Shutdown** check boxes. The **Sysprep** tool makes the guest unique by generating an OS ID, which is used by certain Microsoft services.

> **IMPORTANT**
>
> Red Hat does not provide technical support for Cloudbase-Init. If you encounter an issue, contact Cloudbase Solutions .

When the Windows system shuts down, the *name*.**qcow2** image file is ready to be uploaded to the Image service. For more information on uploading this image to your OpenStack deployment using the dashboard or the command line, see Upload an Image.

1.2.1.3. Use libosinfo

Image Service (glance) can process *libosinfo* data for images, making it easier to configure the optimal virtual hardware for an instance. This can be done by adding the libosinfo-formatted operating system name to the glance image.

1. This example specifies that the image with ID **654dbfd5-5c01-411f-8599-a27bd344d79b** uses the libosinfo value of **rhel7.2**:

   ```
   $ openstack image set 654dbfd5-5c01-411f-8599-a27bd344d79b --property
   os_name=rhel7.2
   ```

   As a result, Compute will supply virtual hardware optimized for **rhel7.2** whenever an instance is built using the **654dbfd5-5c01-411f-8599-a27bd344d79b** image.

> **NOTE**
>
> For a complete list of **libosinfo** values, refer to the libosinfo project:
> https://gitlab.com/libosinfo/osinfo-db/tree/master/data/os

1.2.2. Upload an Image

1. In the dashboard, select **Project > Compute > Images**

2. Click **Create Image**.

3. Fill out the values, and click **Create Image** when finished.

Table 1.1. Image Options

Field	Notes
Name	Name for the image. The name must be unique within the project.
Description	Brief description to identify the image.
Image Source	Image source: **Image Location** or **Image File**. Based on your selection, the next field is displayed.
Image Location or Image File	- Select **Image Location** option to specify the image location URL. - Select **Image File** option to upload an image from the local disk.
Format	Image format (for example, qcow2).
Architecture	Image architecture. For example, use i686 for a 32-bit architecture or x86_64 for a 64-bit architecture.
Minimum Disk (GB)	Minimum disk size required to boot the image. If this field is not specified, the default value is 0 (no minimum).
Minimum RAM (MB)	Minimum memory size required to boot the image. If this field is not specified, the default value is 0 (no minimum).
Public	If selected, makes the image public to all users with access to the project.
Protected	If selected, ensures only users with specific permissions can delete this image.

When the image has been successfully uploaded, its status is changed to **active**, which indicates that the image is available for use. Note that the Image service can handle even large images that take a long time to upload — longer than the lifetime of the Identity service token which was used when the upload was initiated. This is due to the fact that the Image service first creates a trust with the Identity service so that a new token can be obtained and used when the upload is complete and the status of the image is to be updated.

> **NOTE**
>
> You can also use the **glance image-create** command with the **property** option to upload an image. More values are available on the command line. For a complete listing, see Image Configuration Parameters.

1.2.3. Update an Image

1. In the dashboard, select **Project > Compute > Images**

2. Click **Edit Image** from the dropdown list.

> **NOTE**
>
> The **Edit Image** option is available only when you log in as an **admin** user. When you log in as a **demo** user, you have the option to **Launch an instance** or **Create Volume**.

3. Update the fields and click **Update Image** when finished. You can update the following values – name, description, kernel ID, ramdisk ID, architecture, format, minimum disk, minimum RAM, public, protected.

4. Click the drop-down menu and select **Update Metadata** option.

5. Specify metadata by adding items from the left column to the right one. In the left column, there are metadata definitions from the Image Service Metadata Catalog. Select **Other** to add metadata with the key of your choice and click **Save** when finished.

> **NOTE**
>
> You can also use the **glance image-update** command with the **property** option to update an image. More values are available on the command line; for a complete listing, see Image Configuration Parameters.

1.2.4. Import an Image

You can import images into the Image service (glance) using **web-download** to import an image from a URI and **glance-direct** to import an image from a local file system. Both options are enabled by default.

Import methods are configured by the cloud administrator. Run the **glance import-info** command to list available import options.

1.2.4.1. Import from a Remote URI

You can use the **web-download** method to copy an image from a remote URI.

1. Create an image and specify the URI of the image to import.

   ```
   glance image-create --uri <URI>
   ```

2. You can monitor the image's availability using the **glance image-show <image-ID>** command where the ID is the one provided during image creation.

The Image service web-download method uses a two-stage process to perform the import. First, it creates an image record. Second, it retrieves the image the specified URI. This method provides a more secure way to import images than the deprecated **copy-from** method used in Image API v1.

The URI is subject to optional blacklist and whitelist filtering as described in the Advanced Overcloud Customization Guide.

The Image Property Injection plugin may inject metadata properties to the image as described in the Advanced Overcloud Customization Guide. These injected properties determine which compute nodes the image instances are launched on.

1.2.4.2. Import from a Local Volume

The **glance-direct** method creates an image record, which generates an image ID. Once the image is

uploaded to the service from a local volume, it is stored in a staging area and is made active after it passes any configured checks. The **glance-direct** method requires a shared staging area when used in a highly available (HA) configuration.

> **NOTE**
>
> Image uploads using the **glance-direct** method fail in an HA environment if a common staging area is not present. In an HA active–active environment, API calls are distributed to the glance controllers. The download API call could be sent to a different controller than the API call to upload the image. For more information about configuring the staging area, refer to the Storage Configuration section in the *Advanced Overcloud Customization Guide*.

The glance-direct method uses three different calls to import an image:

- **glance image-create**

- **glance image-stage**

- **glance image-import**

You can use the **glance image-create-via-import** command to perform all three of these calls in one command. In the example below, uppercase words should be replaced with the appropriate options.

```
glance image-create-via-import --container-format FORMAT --disk-format DISKFORMAT --name
NAME --file /PATH/TO/IMAGE
```

Once the image moves from the staging area to the back end location, the image is listed. However, it may take some time for the image to become active.

You can monitor the image's availability using the **glance image-show <image-ID>** command where the ID is the one provided during image creation.

1.2.5. Delete an Image

1. In the dashboard, select **Project > Compute > Images**

2. Select the image you want to delete and click **Delete Images**.

1.2.6. Hide or Unhide an Image

You can hide public images from normal listings presented to users. For instance, you can hide obsolete CentOS 7 images and show only the latest version to simplify the user experience. Users can discover and use hidden images.

To hide an image:

```
glance image-update <image-id> --hidden 'true'
```

To create a hidden image, add the **--hidden** argument to the **glance image-create** command.

To unhide an image:

```
glance image-update <image-id> --hidden 'false'
```

1.2.7. Show Hidden Images

To list hidden images:

```
glance image-list --hidden 'true'
```

CHAPTER 2. CONFIGURING THE COMPUTE (NOVA) SERVICE

Use environment files to customize the Compute (nova) service. Puppet generates and stores this configuration in the **/var/lib/config-data/puppet-generated/<nova_container>/etc/nova/nova.conf** file. Use the following configuration methods to customize the Compute service configuration:

- **Heat parameters** - as detailed in the Compute (nova) Parameters section in the *Overcloud Parameters* guide. For example:

```
parameter_defaults:
  NovaSchedulerDefaultFilters:
AggregateInstanceExtraSpecsFilter,RetryFilter,ComputeFilter,ComputeCapabilitiesFilter,Image
PropertiesFilter
  NovaNfsEnabled: true
  NovaNfsShare: '192.0.2.254:/export/nova'
  NovaNfsOptions: 'context=system_u:object_r:nfs_t:s0'
  NovaNfsVersion: '4.2'
```

- **Puppet parameters** - as defined in **/etc/puppet/modules/nova/manifests/***:

```
parameter_defaults:
  ComputeExtraConfig:
    nova::compute::force_raw_images: True
```

> **NOTE**
>
> Only use this method if an equivalent heat parameter does not exist.

- **Manual hieradata overrides** - for customizing parameters when no heat or Puppet parameter exists. For example, the following sets the **disk_allocation_ratio** in the **[DEFAULT]** section on the Compute role:

```
parameter_defaults:
  ComputeExtraConfig:
    nova::config::nova_config:
      DEFAULT/disk_allocation_ratio:
        value: '2.0'
```

> **WARNING**
>
> If a heat parameter exists, it must be used instead of the Puppet parameter; if a Puppet parameter exists, but not a heat parameter, then the Puppet parameter must be used instead of the manual override method. The manual override method must only be used if there is no equivalent heat or Puppet parameter.

TIP

Follow the guidance in Identifying Parameters to Modify to determine if a heat or Puppet parameter is available for customizing a particular configuration.

See Parameters in the *Advanced Overcloud Customization* guide for further details on configuring overcloud services.

CHAPTER 3. CONFIGURE OPENSTACK COMPUTE STORAGE

This chapter describes the architecture for the back-end storage of images in OpenStack Compute (nova), and provides basic configuration options.

3.1. ARCHITECTURE OVERVIEW

In Red Hat OpenStack Platform, the OpenStack Compute service uses the KVM hypervisor to execute compute workloads. The **libvirt** driver handles all interactions with KVM, and enables the creation of virtual machines.

Two types of **libvirt** storage must be considered for Compute:

- Base image, which is a cached and formatted copy of the Image service image.

- Instance disk, which is created using the **libvirt** base and is the back end for the virtual machine instance. Instance disk data can be stored either in Compute's ephemeral storage (using the **libvirt** base) or in persistent storage (for example, using Block Storage).
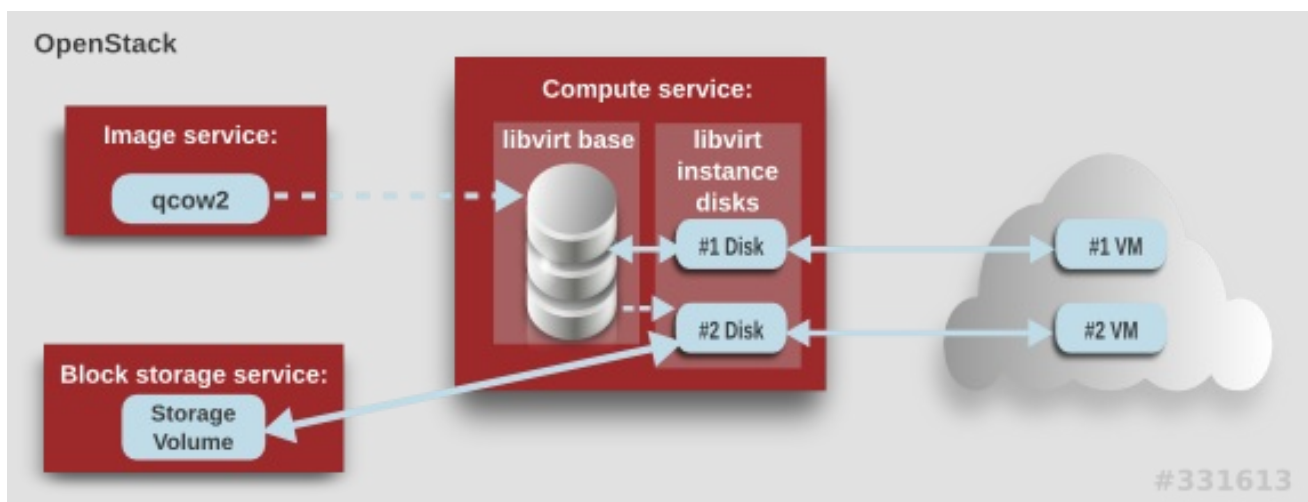
The steps that Compute takes to create a virtual machine instance are:

1. Cache the Image service's backing image as the **libvirt** base.

2. Convert the base image to the raw format (if configured).

3. Resize the base image to match the VM's flavor specifications.

4. Use the base image to create the libvirt instance disk.

In the diagram above, the #1 instance disk uses ephemeral storage; the #2 disk uses a block-storage volume.

Ephemeral storage is an empty, unformatted, additional disk available to an instance. This storage value is defined by the instance flavor. The value provided by the user must be less than or equal to the ephemeral value defined for the flavor. The default value is **0**, meaning no ephemeral storage is created.

The ephemeral disk appears in the same way as a plugged-in hard drive or thumb drive. It is available as a block device which you can check using the **lsblk** command. You can format it, mount it, and use it however you normally would a block device. There is no way to preserve or reference that disk beyond the instance it is attached to.

Block storage volume is persistant storage available to an instance regardless of the state of the running instance.

3.2. CONFIGURATION

You can configure performance tuning and security for your virtual disks by customizing the Compute (nova) configuration files. Compute is configured in custom environment files and heat templates using the parameters detailed in the Compute (nova) Parameters section in the *Overcloud Parameters* guide. This configuration is generated and stored in the **/var/lib/config-data/puppet-generated/<nova_container>/etc/nova/nova.conf** file, as detailed in the following table.

Table 3.1. Compute Image Parameters

Section	Parameter	Description	Default
[DEFAULT]	**force_raw_images**	Whether to convert a **non-raw** cached base image to be **raw** (boolean). If a non-raw image is converted to raw, Compute: ● Disallows backing files (which might be a security issue). ● Removes existing compression (to avoid CPU bottlenecks). Converting the base to raw uses more space for any image that could have been used directly by the hypervisor (for example, a qcow2 image). If you have a system with slower I/O or less available space, you might want to specify *false*, trading the higher CPU requirements of compression for that of minimized input bandwidth. Raw base images are always used with **libvirt_images_type=lvm**.	**true**
[DEFAULT]	**use_cow_images**	Whether to use CoW (Copy on Write) images for **libvirt** instance disks (boolean): ● **false** - The raw format is used. Without CoW, more space is used for common parts of the disk image ● **true** - The cqow2 format is used. With CoW, depending on the backing store and host caching, there may be better concurrency achieved by having each VM operate on its own copy.	true

Section	Parameter	Description	Default
[DEFAULT]	**preallocate_i mages**	Preallocation mode for **libvirt** instance disks. Value can be: • **none** - No storage is provisioned at instance start. • **space** - Storage is fully allocated at instance start (using **fallocate**), which can help with both space guarantees and I/O performance. Even when not using CoW instance disks, the copy each VM gets is sparse and so the VM may fail unexpectedly at run time with ENOSPC. By running **fallocate(1)** on the instance disk images, Compute immediately and efficiently allocates the space for them in the file system (if supported). Run time performance should also be improved because the file system does not have to dynamically allocate blocks at run time (reducing CPU overhead and more importantly file fragmentation).	none
[DEFAULT]	**resize_fs_usin g_block_devi ce**	Whether to enable direct resizing of the base image by accessing the image over a block device (boolean). This is only necessary for images with older versions of **cloud-init** (that cannot resize themselves). Because this parameter enables the direct mounting of images which might otherwise be disabled for security reasons, it is not enabled by default.	**false**
[DEFAULT]	**default_ephe meral_format**	The default format that is used for a new ephemeral volume. Value can be: **ext2**, **ext3**, or **ext4**. The **ext4** format provides much faster initialization times than **ext3** for new, large disks. You can also override per instance using the **guest_format** configuration option.	**ext4**
[DEFAULT]	**image_cache_ manager_inte rval**	Number of seconds to wait between runs of the image cache manager, which impacts base caching on libvirt compute nodes. This period is used in the auto removal of unused cached images (see **remove_unused_base_images** and **remove_unused_original_minimum_age_ seconds**).	**2400**

Section	Parameter	Description	Default
[DEFAULT]	**remove_unused_base_images**	Whether to enable the automatic removal of unused base images (checked every **image_cache_manager_interval** seconds). Images are defined as **unused** if they have not been accessed in **remove_unused_original_minimum_age_seconds** seconds.	**true**
[DEFAULT]	**remove_unused_original_minimum_age_seconds**	How old an unused base image must be before being removed from the **libvirt** cache (see **remove_unused_base_images**).	**86400**
[**libvirt**]	**images_type**	Image type to use for **libvirt** instance disks (deprecates **use_cow_images**). Value can be: **raw**, **qcow2**, **lvm**, **rbd**, or **default**. If **default** is specified, the value used for the **use_cow_images** parameter is used.	**default**

CHAPTER 4. VIRTUAL MACHINE INSTANCES

OpenStack Compute is the central component that provides virtual machines on demand. Compute interacts with the Identity service for authentication, Image service for images (used to launch instances), and the dashboard service for the user and administrative interface.

Red Hat OpenStack Platform allows you to easily manage virtual machine instances in the cloud. The Compute service creates, schedules, and manages instances, and exposes this functionality to other OpenStack components. This chapter discusses these procedures along with procedures to add components like key pairs, security groups, host aggregates and flavors. The term *instance* is used by OpenStack to mean a virtual machine instance.

4.1. MANAGE INSTANCES

Before you can create an instance, you need to ensure certain other OpenStack components (for example, a network, key pair and an image or a volume as the boot source) are available for the instance.

This section discusses the procedures to add these components, create and manage an instance. Managing an instance refers to updating, and logging in to an instance, viewing how the instances are being used, resizing or deleting them.

4.1.1. Add Components

Use the following sections to create a network, key pair and upload an image or volume source. These components are used in the creation of an instance and are not available by default. You will also need to create a new security group to allow SSH access to the user.

1. In the dashboard, select **Project**.

2. Select **Network > Networks**, and ensure there is a private network to which you can attach the new instance (to create a network, see Create a Network section in the *Networking Guide*).

3. Select **Compute > Access & Security > Key Pairs**, and ensure there is a key pair (to create a key pair, see Section 4.2.1.1, "Create a Key Pair").

4. Ensure that you have either an image or a volume that can be used as a boot source:

   - To view boot-source images, select the **Images** tab (to create an image, see Section 1.2.1, "Create an Image").

   - To view boot-source volumes, select the **Volumes** tab (to create a volume, see Create a Volume in the *Storage Guide*).

5. Select **Compute > Access & Security > Security Groups**, and ensure you have created a security group rule (to create a security group, see Project Security Management in the *Users and Identity Management Guide*).

4.1.2. Launch an Instance

Launch one or more instances from the dashboard.

> **NOTE**
>
> By default, the Launch Instance form is used to launch instances. However, you can also enable a Launch Instance wizard that simplifies the steps required. For more information, see Appendix B, *Enabling the Launch Instance Wizard* .

1. In the dashboard, select **Project > Compute > Instances**

2. Click **Launch Instance**.

3. Fill out the fields (those marked with '* ' are required), and click **Launch**.

One or more instances are created, and launched based on the options provided.

4.1.2.1. Launch Instance Options

The following table outlines the options available when launching a new instance using the Launch Instance form. The same options are also available in the Launch instance wizard.

Table 4.1. Launch Instance Form Options

Tab	Field	Notes
Project and User	Project	Select the project from the dropdown list.
	User	Select the user from the dropdown list.
Details	Availability Zone	Zones are logical groupings of cloud resources in which your instance can be placed. If you are unsure, use the default zone (for more information, see Section 4.4, "Manage Host Aggregates").
	Instance Name	A name to identify your instance.
	Flavor	The flavor determines what resources the instance is given (for example, memory). For default flavor allocations and information on creating new flavors, see Section 4.3, "Manage Flavors".
	Instance Count	The number of instances to create with these parameters. "1" is preselected.
	Instance Boot Source	Depending on the item selected, new fields are displayed allowing you to select the source: • Image sources must be compatible with OpenStack (see Section 1.2, "Manage Images"). • If a volume or volume source is selected, the source must be formatted using an image (see Basic Volume Usage and Configuration in the *Storage Guide*).

Tab	Field	Notes
Access and Security	Key Pair	The specified key pair is injected into the instance and is used to remotely access the instance using SSH (if neither a direct login information or a static key pair is provided). Usually one key pair per project is created.
	Security Groups	Security groups contain firewall rules which filter the type and direction of the instance's network traffic (for more information on configuring groups, see Project Security Management in the *Users and Identity Management Guide*).
Networking	Selected Networks	You must select at least one network. Instances are typically assigned to a private network, and then later given a floating IP address to enable external access.
Post-Creation	Customization Script Source	You can provide either a set of commands or a script file, which will run after the instance is booted (for example, to set the instance host name or a user password). If *Direct Input* is selected, write your commands in the Script Data field; otherwise, specify your script file. **NOTE** Any script that starts with *#cloud-config* is interpreted as using the cloud-config syntax (for information on the syntax, see http://cloudinit.readthedocs.org/en/latest/topics/examples.html).
Advanced Options	Disk Partition	By default, the instance is built as a single partition and dynamically resized as needed. However, you can choose to manually configure the partitions yourself.
	Configuration Drive	If selected, OpenStack writes metadata to a read-only configuration drive that is attached to the instance when it boots (instead of to Compute's metadata service). After the instance has booted, you can mount this drive to view its contents (enables you to provide files to the instance).

4.1.3. Update an Instance (Actions menu)

You can update an instance by selecting **Project > Compute > Instances** and selecting an action for that instance in the **Actions** column. Actions allow you to manipulate the instance in a number of ways:

Table 4.2. Update Instance Options

Action	Description

Action	Description
Create Snapshot	Snapshots preserve the disk state of a running instance. You can create a snapshot to migrate the instance, as well as to preserve backup copies.
Associate/Disassociate Floating IP	You must associate an instance with a floating IP (external) address before it can communicate with external networks, or be reached by external users. Because there are a limited number of external addresses in your external subnets, it is recommended that you disassociate any unused addresses.
Edit Instance	Update the instance's name and associated security groups.
Edit Security Groups	Add and remove security groups to or from this instance using the list of available security groups (for more information on configuring groups, see Project Security Management in the *Users and Identity Management Guide*).
Console	View the instance's console in the browser (allows easy access to the instance).
View Log	View the most recent section of the instance's console log. Once opened, you can view the full log by clicking View Full Log.
Pause/Resume Instance	Immediately pause the instance (you are not asked for confirmation); the state of the instance is stored in memory (RAM).
Suspend/Resume Instance	Immediately suspend the instance (you are not asked for confirmation); like hibernation, the state of the instance is kept on disk.
Resize Instance	Bring up the Resize Instance window (see Section 4.1.4, "Resize an Instance").
Soft Reboot	Gracefully stop and restart the instance. A soft reboot attempts to gracefully shut down all processes before restarting the instance.
Hard Reboot	Stop and restart the instance. A hard reboot effectively just shuts down the instance's *power* and then turns it back on.
Shut Off Instance	Gracefully stop the instance.

Action	Description
Rebuild Instance	Use new image and disk-partition options to rebuild the image (shut down, re-image, and re-boot the instance). If encountering operating system issues, this option is easier to try than terminating the instance and starting over.
Terminate Instance	Permanently destroy the instance (you are asked for confirmation).

You can create and allocate an external IP address, see Section 4.2.3, "Create, Assign, and Release Floating IP Addresses"

4.1.4. Resize an Instance

To resize an instance (memory or CPU count), you must select a new flavor for the instance that has the right capacity. If you are increasing the size, remember to first ensure that the host has enough space.

1. Ensure communication between hosts by setting up each host with SSH key authentication so that Compute can use SSH to move disks to other hosts (for example, compute nodes can share the same SSH key).

2. Enable resizing on the original host by setting the **allow_resize_to_same_host** parameter to "True" in your Compute environment file.

The **allow_resize_to_same_host** parameter does not resize the instance on the same host. Even if the parameter equals **True** on all Compute nodes, the scheduler does not force the instance to resize on the same host. This is the expected behavior.

1. In the dashboard, select **Project > Compute > Instances**

2. Click the instance's **Actions** arrow, and select **Resize Instance**.

3. Select a new flavor in the **New Flavor** field.

4. If you want to manually partition the instance when it launches (results in a faster build time):

   a. Select **Advanced Options**.

   b. In the **Disk Partition** field, select **Manual**.

5. Click **Resize**.

4.1.5. Connect to an Instance

This section discusses the different methods you can use to access an instance console using the dashboard or the command-line interface. You can also directly connect to an instance's serial port allowing you to debug even if the network connection fails.

4.1.5.1. Access an Instance Console using the Dashboard

The console allows you a way to directly access your instance within the dashboard.

1. In the dashboard, select **Compute > Instances**.

2. Click the instance's **More** button and select **Console**.
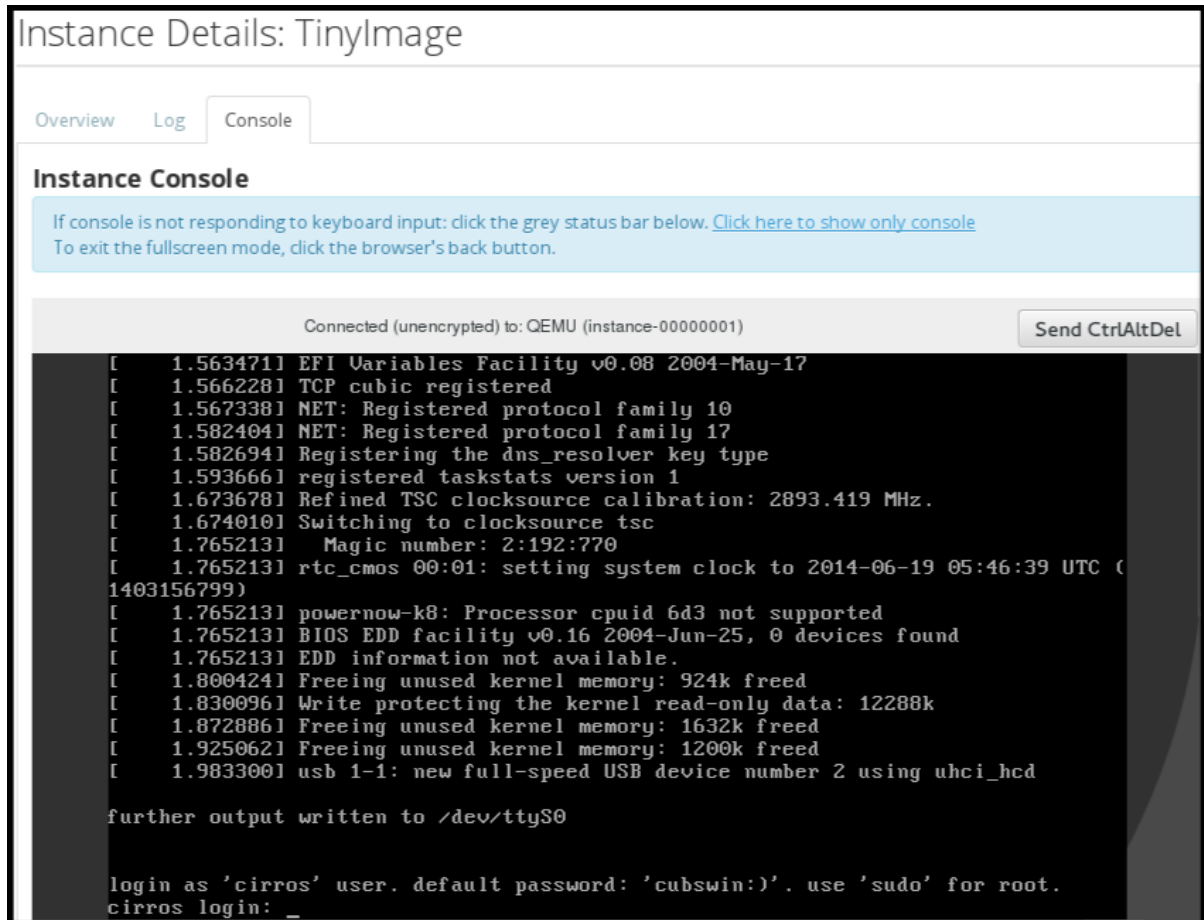
3. Log in using the image's user name and password (for example, a CirrOS image uses *cirros/cubswin:)*).

4.1.5.2. Directly Connect to a VNC Console

You can directly access an instance's VNC console using a URL returned by **nova get-vnc-console** command.

Browser

To obtain a browser URL, use:

```
$ nova get-vnc-console INSTANCE_ID novnc
```

Java Client

To obtain a Java-client URL, use:

```
$ nova get-vnc-console INSTANCE_ID xvpvnc
```

> **NOTE**
>
> nova-xvpvncviewer provides a simple example of a Java client. To download the client, use:
>
> ```
> # git clone https://github.com/cloudbuilders/nova-xvpvncviewer
> # cd nova-xvpvncviewer/viewer
> # make
> ```
>
> Run the viewer with the instance's Java-client URL:
>
> ```
> # java -jar VncViewer.jar URL
> ```
>
> This tool is provided only for customer convenience, and is not officially supported by Red Hat.

4.1.6. View Instance Usage

The following usage statistics are available:

- Per Project
  To view instance usage per project, select **Project > Compute > Overview**. A usage summary is immediately displayed for all project instances.

  You can also view statistics for a specific period of time by specifying the date range and clicking **Submit**.

- Per Hypervisor
  If logged in as an administrator, you can also view information for all projects. Click **Admin > System** and select one of the tabs. For example, the **Resource Usage** tab offers a way to view reports for a distinct time period. You might also click **Hypervisors** to view your current vCPU, memory, or disk statistics.

  > **NOTE**
  >
  > The **vCPU Usage** value (**x of y**) reflects the number of total vCPUs of all virtual machines (x) and the total number of hypervisor cores (y).

4.1.7. Delete an Instance

1. In the dashboard, select **Project > Compute > Instances** and select your instance.

2. Click **Terminate Instance**.

> **NOTE**
>
> Deleting an instance does not delete its attached volumes; you must do this separately (see Delete a Volume in the *Storage Guide*).

4.1.8. Manage Multiple Instances at Once

If you need to start multiple instances at the same time (for example, those that were down for compute or controller maintenance) you can do so easily at **Project > Compute > Instances**

1. Click the check boxes in the first column for the instances that you want to start. If you want to select all of the instances, click the check box in the first row in the table.

2. Click **More Actions** above the table and select **Start Instances**.

Similarly, you can shut off or soft reboot multiple instances by selecting the respective actions.

4.2. MANAGE INSTANCE SECURITY

You can manage access to an instance by assigning it the correct security group (set of firewall rules) and key pair (enables SSH user access). Further, you can assign a floating IP address to an instance to enable external network access. The sections below outline how to create and manage key pairs, security groups, floating IP addresses and logging in to an instance using SSH. There is also a procedure for injecting an **admin** password in to an instance.

For information on managing security groups, see Project Security Management in the *Users and Identity Management Guide*.

4.2.1. Manage Key Pairs

Key pairs provide SSH access to the instances. Each time a key pair is generated, its certificate is downloaded to the local machine and can be distributed to users. Typically, one key pair is created for each project (and used for multiple instances).

You can also import an existing key pair into OpenStack.

4.2.1.1. Create a Key Pair

1. In the dashboard, select **Project > Compute > Access & Security**

2. On the **Key Pairs** tab, click **Create Key Pair**.

3. Specify a name in the **Key Pair Name** field, and click **Create Key Pair**.

When the key pair is created, a key pair file is automatically downloaded through the browser. Save this file for later connections from external machines. For command-line SSH connections, you can load this file into SSH by executing:

```
# ssh-add ~/.ssh/os-key.pem
```

4.2.1.2. Import a Key Pair

1. In the dashboard, select **Project > Compute > Access & Security**

2. On the **Key Pairs** tab, click **Import Key Pair**.

3. Specify a name in the **Key Pair Name** field, and copy and paste the contents of your public key into the **Public Key** field.

4. Click **Import Key Pair**.

4.2.1.3. Delete a Key Pair

1. In the dashboard, select **Project > Compute > Access & Security**

2. On the **Key Pairs** tab, click the key's **Delete Key Pair** button.

4.2.2. Create a Security Group

Security groups are sets of IP filter rules that can be assigned to project instances, and which define networking access to the instance. Security group are project specific; project members can edit the default rules for their security group and add new rule sets.

1. In the dashboard, select the **Project** tab, and click **Compute > Access & Security**.

2. On the **Security Groups** tab, click **+ Create Security Group**.

3. Provide a name and description for the group, and click **Create Security Group**.

For more information on managing project security, see Project Security Management in the *Users and Identity Management Guide*.

4.2.3. Create, Assign, and Release Floating IP Addresses

By default, an instance is given an internal IP address when it is first created. However, you can enable access through the public network by creating and assigning a floating IP address (external address). You can change an instance's associated IP address regardless of the instance's state.

Projects have a limited range of floating IP address that can be used (by default, the limit is 50), so you should release these addresses for reuse when they are no longer needed. Floating IP addresses can only be allocated from an existing floating IP pool, see Create Floating IP Pools in the *Networking Guide*.

4.2.3.1. Allocate a Floating IP to the Project

1. In the dashboard, select **Project > Compute > Access & Security**

2. On the **Floating IPs** tab, click **Allocate IP to Project**.

3. Select a network from which to allocate the IP address in the **Pool** field.

4. Click **Allocate IP**.

4.2.3.2. Assign a Floating IP

1. In the dashboard, select **Project > Compute > Access & Security**

2. On the **Floating IPs** tab, click the address' **Associate** button.

3. Select the address to be assigned in the IP address field.

> **NOTE**
>
> If no addresses are available, you can click the **+** button to create a new address.

4. Select the instance to be associated in the **Port** to be **Associated** field. An instance can only be associated with one floating IP address.

5. Click **Associate**.

4.2.3.3. Release a Floating IP

1. In the dashboard, select **Project > Compute > Access & Security**

2. On the **Floating IPs** tab, click the address' menu arrow (next to the **Associate/Disassociate** button).

3. Select **Release Floating IP**.

4.2.4. Log in to an Instance

Prerequisites:

- Ensure that the instance's security group has an SSH rule (see Project Security Management in the *Users and Identity Management Guide*).

- Ensure the instance has a floating IP address (external address) assigned to it (see Section 4.2.3, "Create, Assign, and Release Floating IP Addresses").

- Obtain the instance's key-pair certificate. The certificate is downloaded when the key pair is created; if you did not create the key pair yourself, ask your administrator (see Section 4.2.1, "Manage Key Pairs").

To first load the key pair file into SSH, and then use ssh without naming it

1. Change the permissions of the generated key-pair certificate.

   ```
   $ chmod 600 os-key.pem
   ```

2. Check whether **ssh-agent** is already running:

   ```
   # ps -ef | grep ssh-agent
   ```

3. If not already running, start it up with:

   ```
   # eval `ssh-agent`
   ```

4. On your local machine, load the key-pair certificate into SSH. For example:

   ```
   $ ssh-add ~/.ssh/os-key.pem
   ```

5. You can now SSH into the file with the user supplied by the image.

The following example command shows how to SSH into the Red Hat Enterprise Linux guest image with the user **cloud-user**:

```
$ ssh cloud-user@192.0.2.24
```

> **NOTE**
>
> You can also use the certificate directly. For example:
>
> ```
> $ ssh -i /myDir/os-key.pem cloud-user@192.0.2.24
> ```

4.2.5. Inject an `admin` Password Into an Instance

You can inject an **admin** (**root**) password into an instance using the following procedure.

1. In the **/etc/openstack-dashboard/local_settings** file, set the **change_set_password** parameter value to **True**.

   ```
   can_set_password: True
   ```

2. Set the **inject_password** parameter to "True" in your Compute environment file.

   ```
   inject_password=true
   ```

3. Restart the Compute service.

   ```
   # service nova-compute restart
   ```

When you use the **nova boot** command to launch a new instance, the output of the command displays an **adminPass** parameter. You can use this password to log into the instance as the   **root** user.

The Compute service overwrites the password value in the **/etc/shadow** file for the **root** user. This procedure can also be used to activate the **root** account for the KVM guest images. For more information on how to use KVM guest images, see Section 1.2.1.1, "Use a KVM Guest Image With Red Hat OpenStack Platform"

You can also set a custom password from the dashboard. To enable this, run the following command after you have set **can_set_password** parameter to **true**.

```
# systemctl restart httpd.service
```

The newly added **admin** password fields are as follows:

These fields can be used when you launch or rebuild an instance.

4.3. MANAGE FLAVORS

Each created instance is given a flavor (resource template), which determines the instance's size and capacity. Flavors can also specify secondary ephemeral storage, swap disk, metadata to restrict usage, or special project access (none of the default flavors have these additional attributes defined).

Table 4.3. Default Flavors

Name	vCPUs	RAM	Root Disk Size
m1.tiny	1	512 MB	1 GB
m1.small	1	2048 MB	20 GB
m1.medium	2	4096 MB	40 GB
m1.large	4	8192 MB	80 GB
m1.xlarge	8	16384 MB	160 GB

The majority of end users will be able to use the default flavors. However, you can create and manage specialized flavors. For example, you can:

- Change default memory and capacity to suit the underlying hardware needs.

- Add metadata to force a specific I/O rate for the instance or to match a host aggregate.

> **NOTE**
>
> Behavior set using image properties overrides behavior set using flavors (for more information, see Section 1.2, "Manage Images").

4.3.1. Update Configuration Permissions

By default, only administrators can create flavors or view the complete flavor list (select Admin > System > Flavors). To allow all users to configure flavors, specify the following in the **/etc/nova/policy.json** file (nova-api server):

"compute_extension:flavormanage": "",

4.3.2. Create a Flavor

1. As an admin user in the dashboard, select **Admin > System > Flavors**

2. Click **Create Flavor**, and specify the following fields:

Table 4.4. Flavor Options

Tab	Field	Description
Flavor Information	Name	Unique name.
	ID	Unique ID. The default value, **auto**, generates a UUID4 value, but you can also manually specify an integer or UUID4 value.
	VCPUs	Number of virtual CPUs.
	RAM (MB)	Memory (in megabytes).
	Root Disk (GB)	Ephemeral disk size (in gigabytes); to use the native image size, specify **0**. This disk is not used if **Instance Boot Source=Boot from Volume**.
	Epehemeral Disk (GB)	Secondary ephemeral disk size (in gigabytes) available to an instance. This disk is destroyed when an instance is deleted. The default value is **0**, which implies that no ephemeral disk is created.

Tab	Field	Description
	Swap Disk (MB)	Swap disk size (in megabytes).
Flavor Access	Selected Projects	Projects which can use the flavor. If no projects are selected, all projects have access (**Public=Yes**).

3. Click Create Flavor.

4.3.3. Update General Attributes

1. As an admin user in the dashboard, select **Admin > System > Flavors**

2. Click the flavor's **Edit Flavor** button.

3. Update the values, and click **Save**.

4.3.4. Update Flavor Metadata

In addition to editing general attributes, you can add metadata to a flavor (**extra_specs**), which can help fine-tune instance usage. For example, you might want to set the maximum-allowed bandwidth or disk writes.

- Pre-defined keys determine hardware support or quotas. Pre-defined keys are limited by the hypervisor you are using (for libvirt, see Table 4.5, "Libvirt Metadata").

- Both pre-defined and user-defined keys can determine instance scheduling. For example, you might specify **SpecialComp=True**; any instance with this flavor can then only run in a host aggregate with the same key-value combination in its metadata (see Section 4.4, "Manage Host Aggregates").

4.3.4.1. View Metadata

1. As an admin user in the dashboard, select **Admin > System > Flavors**

2. Click the flavor's **Metadata** link (**Yes** or **No**). All current values are listed on the right-hand side under **Existing Metadata**.

4.3.4.2. Add Metadata

You specify a flavor's metadata using a **key**/**value** pair.

1. As an admin user in the dashboard, select **Admin > System > Flavors**

2. Click the flavor's **Metadata** link (**Yes** or **No**). All current values are listed on the right-hand side under **Existing Metadata**.

3. Under **Available Metadata**, click on the **Other** field, and specify the key you want to add (see Table 4.5, "Libvirt Metadata").

4. Click the + button; you can now view the new key under **Existing Metadata**.

5. Fill in the key's value in its right-hand field.

6. When finished with adding key-value pairs, click **Save**.

Table 4.5. Libvirt Metadata

Key	Description
hw:action	Action that configures support limits per instance. Valid actions are: • **cpu_max_sockets** - Maximum supported CPU sockets. • **cpu_max_cores** - Maximum supported CPU cores. • **cpu_max_threads** - Maximum supported CPU threads. • **cpu_sockets** - Preferred number of CPU sockets. • **cpu_cores** - Preferred number of CPU cores. • **cpu_threads** - Preferred number of CPU threads. • **serial_port_count** - Maximum serial ports per instance. Example: **hw:cpu_max_sockets=2**

Key	Description
hw:NUMA_def	Definition of NUMA topology for the instance. For flavors whose RAM and vCPU allocations are larger than the size of NUMA nodes in the compute hosts, defining NUMA topology enables hosts to better utilize NUMA and improve performance of the guest OS. NUMA definitions defined through the flavor override image definitions. Valid definitions are:

- **numa_nodes** - Number of NUMA nodes to expose to the instance. Specify *1* to ensure image NUMA settings are overridden.

- **numa_cpus.0** - Mapping of vCPUs N-M to NUMA node 0 (comma-separated list).

- **numa_cpus.1** - Mapping of vCPUs N-M to NUMA node 1 (comma-separated list).

- **numa_mem.0** - Mapping N MB of RAM to NUMA node 0.

- **numa_mem.1** - Mapping N MB of RAM to NUMA node 1.

- **numa_cpu.N** and **numa_mem.N** are only valid if **numa_nodes** is set. Additionally, they are only required if the instance's NUMA nodes have an asymetrical allocation of CPUs and RAM (important for some NFV workloads).

> **NOTE**
>
> If the values of **numa_cpu** or **numa_mem.N** specify more than that available, an exception is raised.

Example when the instance has 8 vCPUs and 4GB RAM:

- **hw:numa_nodes=2**

- **hw:numa_cpus.0=0,1,2,3,4,5**

- **hw:numa_cpus.1=6,7**

- **hw:numa_mem.0=3072**

- **hw:numa_mem.1=1024**

The scheduler looks for a host with 2 NUMA nodes with the ability to run 6 CPUs + 3072 MB, or 3 GB, of RAM on one node, and 2 CPUS + 1024 MB, or 1 GB, of RAM on another node. If a host has a single NUMA node with capability to run 8 CPUs and 4 GB of RAM, it will not be considered a valid match.

Key	Description
hw:watchdog_action	An instance watchdog device can be used to trigger an action if the instance somehow fails (or hangs). Valid actions are: • **disabled** – The device is not attached (default value). • **pause** – Pause the instance. • **poweroff** – Forcefully shut down the instance. • **reset** – Forcefully reset the instance. • **none** – Enable the watchdog, but do nothing if the instance fails. Example: **hw:watchdog_action=poweroff**
hw_rng:action	A random-number generator device can be added to an instance using its image properties (see **hw_rng_model** in the "Command-Line Interface Reference" in Red Hat OpenStack Platform documentation). If the device has been added, valid actions are: • **allowed** – If **True**, the device is enabled; if **False**, disabled. By default, the device is disabled. • **rate_bytes** – Maximum number of bytes the instance's kernel can read from the host to fill its entropy pool every rate_period (integer). • **rate_period** – Duration of the read period in seconds (integer). Example: **hw_rng:allowed=True**.
hw_video:ram_max_mb	Maximum permitted RAM to be allowed for video devices (in MB). Example: **hw:ram_max_mb=64**
quota:option	Enforcing limit for the instance. Valid options are: • **cpu_period** – Time period for enforcing **cpu_quota** (in microseconds). Within the specified **cpu_period**, each vCPU cannot consume more than **cpu_quota** of runtime. The value must be in range [1000, 1000000]; *0* means *no value*. • **cpu_quota - Maximum allowed bandwidth (in microseconds) for the vCPU in each `cpu_period**`. The value must be in range [1000, 18446744073709551]. *0* means *no value*; a negative value means that the vCPU is not controlled. **cpu_quota** and **cpu_period** can be used to ensure that all vCPUs run at the same speed. • **cpu_shares** – Share of CPU time for the domain. The value only has meaning when weighted against other machine values in the same domain. That is, an instance with a flavor with *200* will get twice as much machine time as an instance with *100*. • **disk_read_bytes_sec** – Maximum disk reads in bytes per second. • **disk_read_iops_sec** – Maximum read I/O operations per second.

Key	Description
	• **disk_write_bytes_sec** – Maximum disk writes in bytes per second.

• **disk_write_iops_sec** – Maximum write I/O operations per second.

- **disk_total_bytes_sec** – Maximum total throughput limit in bytes per second.

- **disk_total_iops_sec** – Maximum total I/O operations per second.

- **vif_inbound_average** – Desired average of incoming traffic.

- **vif_inbound_burst** – Maximum amount of traffic that can be received at **vif_inbound_peak** speed.

- **vif_inbound_peak** – Maximum rate at which incoming traffic can be received.

- **vif_outbound_average** – Desired average of outgoing traffic.

- **vif_outbound_burst** – Maximum amount of traffic that can be sent at **vif_outbound_peak** speed.

- **vif_outbound_peak** – Maximum rate at which outgoing traffic can be sent.

Example: **quota:vif_inbound_average=10240**

In addition, the VMware driver supports the following quota options, which control upper and lower limits for CPUs, RAM, disks, and networks, as well as *shares*, which can be used to control relative allocation of available resources among tenants:

- **cpu_limit** – Maximum CPU frequency available to a virtual machine (in MHz).

- **cpu_reservation** – Guaranteed minimum amount of CPU resources available to a virtual machine (in MHz).

- **cpu_shares_level** – CPU allocation level (shares) in the case of contention. Possible values are **high**, **normal**, **low**, and **custom**.

- **cpu_shares_share** – The number of allocated CPU shares. Applicable when **cpu_shares_level** is set to **custom**.

- **memory_limit** – Maximum amount of RAM available to a virtual machine (in MB).

- **memory_reservation** – Guaranteed minimum amount of RAM available to a virtual machine (in MB).

- **memory_shares_level** – RAM allocation level (shares) in the case of contention. Possible values are **high**, **normal**, **low**, and **custom**.

- **memory_shares_share** – The number of allocated RAM shares. Applicable when **memory_shares_level** is set to **custom**.

- **disk_io_limit** – Maximum I/O utilization by a virtual machine (in I/O operations per second).

- **disk_io_reservation** – Guaranteed minimum amount of disk resources available to a virtual machine (in I/O operations per second).

- **disk_io_shares_level** – I/O allocation level (shares) in the case of contention. Possible values are **high**, **normal**, **low**, and **custom**.

Key	Description
	• **disk_io_shares_share** - The number of allocated I/O shares. Applicable when **disk_io_shares_level** is set to **custom**. • **vif_limit** – Maximum network bandwidth available to a virtual network adapter (in Mbps). • **vif_reservation** – Guaranteed minimum network bandwidth available to a virtual network adapter (in Mbps). • **vif_shares_level** – Network bandwidth allocation level (shares) in the case of contention. Possible values are **high**, **normal**, **low**, and **custom**. • **vif_shares_share** – The number of allocated network bandwidth shares. Applicable when **vif_shares_level** is set to **custom**.

4.4. MANAGE HOST AGGREGATES

A single Compute deployment can be partitioned into logical groups for performance or administrative purposes. OpenStack uses the following terms:

- *Host aggregates* – A host aggregate creates logical units in a OpenStack deployment by grouping together hosts. Aggregates are assigned Compute hosts and associated metadata; a host can be in more than one host aggregate. Only administrators can see or create host aggregates.
  An aggregate's metadata is commonly used to provide information for use with the Compute scheduler (for example, limiting specific flavors or images to a subset of hosts). Metadata specified in a host aggregate will limit the use of that host to any instance that has the same metadata specified in its flavor.

  Administrators can use host aggregates to handle load balancing, enforce physical isolation (or redundancy), group servers with common attributes, or separate out classes of hardware. When you create an aggregate, a zone name must be specified, and it is this name which is presented to the end user.

- *Availability zones* – An availability zone is the end-user view of a host aggregate. An end user cannot view which hosts make up the zone, nor see the zone's metadata; the user can only see the zone's name.
  End users can be directed to use specific zones which have been configured with certain capabilities or within certain areas.

4.4.1. Enable Host Aggregate Scheduling

By default, host-aggregate metadata is not used to filter instance usage. You must update the Compute scheduler's configuration to enable metadata usage:

1. Open your Compute environment file.

2. Add the following values to the **NovaSchedulerDefaultFilters** parameter, if they are not already present:

   - **AggregateInstanceExtraSpecsFilter** for host aggregate metadata.

> **NOTE**
>
> Scoped specifications must be used for setting flavor **extra_specs** when specifying both **AggregateInstanceExtraSpecsFilter** and **ComputeCapabilitiesFilter** filters as values of the same **NovaSchedulerDefaultFilters** parameter, otherwise the **ComputeCapabilitiesFilter** will fail to select a suitable host. For details on the namespaces to use to scope the flavor **extra_specs** keys for these filters, see Table 4.7, "Scheduling Filters".

- **AvailabilityZoneFilter** for availability zone host specification when launching an instance.

3. Save the configuration file.

4. Deploy the overcloud.

4.4.2. View Availability Zones or Host Aggregates

As an admin user in the dashboard, select **Admin > System > Host Aggregates** All currently defined aggregates are listed in the **Host Aggregates** section; all zones are in the **Availability Zones** section.

4.4.3. Add a Host Aggregate

1. As an admin user in the dashboard, select **Admin > System > Host Aggregates** All currently defined aggregates are listed in the **Host Aggregates** section.

2. Click **Create Host Aggregate**.

3. Add a name for the aggregate in the **Name** field, and a name by which the end user should see it in the **Availability Zone** field.

4. Click **Manage Hosts within Aggregate**

5. Select a host for use by clicking its + icon.

6. Click **Create Host Aggregate**.

4.4.4. Update a Host Aggregate

1. As an admin user in the dashboard, select **Admin > System > Host Aggregates** All currently defined aggregates are listed in the **Host Aggregates** section.

2. To update the instance's **Name** or **Availability zone**:

   - Click the aggregate's **Edit Host Aggregate** button.

   - Update the **Name** or **Availability Zone** field, and click **Save**.

3. To update the instance's **Assigned hosts**:

   - Click the aggregate's arrow icon under **Actions**.

   - Click **Manage Hosts**.

   - Change a host's assignment by clicking its + or – icon.

- When finished, click **Save**.

4. To update the instance's **Metadata**:

    - Click the aggregate's arrow icon under **Actions**.

    - Click the **Update Metadata** button. All current values are listed on the right-hand side under **Existing Metadata**.

    - Under **Available Metadata**, click on the **Other** field, and specify the key you want to add. Use predefined keys (see Table 4.6, "Host Aggregate Metadata") or add your own (which will only be valid if exactly the same key is set in an instance's flavor).

    - Click the + button; you can now view the new key under **Existing Metadata**.

        > **NOTE**
        >
        > Remove a key by clicking its – icon.

    - Click **Save**.

        Table 4.6. Host Aggregate Metadata

Key	Description
filter_tenant_id	If specified, the aggregate only hosts this tenant (project). Depends on the **AggregateMultiTenancyIsolation** filter being set for the Compute scheduler.

4.4.5. Delete a Host Aggregate

1. As an admin user in the dashboard, select **Admin > System > Host Aggregates** All currently defined aggregates are listed in the **Host Aggregates** section.

2. Remove all assigned hosts from the aggregate:

    a. Click the aggregate's arrow icon under **Actions**.

    b. Click **Manage Hosts**.

    c. Remove all hosts by clicking their – icon.

    d. When finished, click **Save**.

3. Click the aggregate's arrow icon under **Actions**.

4. Click **Delete Host Aggregate** in this and the next dialog screen.

4.5. SCHEDULE HOSTS

The Compute scheduling service determines on which host (or host aggregate), an instance will be placed. As an administrator, you can influence where the scheduler will place an instance. For example, you might want to limit scheduling to hosts in a certain group or with the right RAM.

You can configure the following components:

- Filters – Determine the initial set of hosts on which an instance might be placed (see Section 4.5.1, "Configure Scheduling Filters").

- Weights – When filtering is complete, the resulting set of hosts are prioritized using the weighting system. The highest weight has the highest priority (see Section 4.5.2, "Configure Scheduling Weights").

- Scheduler service – There are a number of configuration options in the **/var/lib/config-data/puppet-generated/<nova_container>/etc/nova/nova.conf** file (on the scheduler host), which determine how the scheduler executes its tasks, and handles weights and filters.

In the following diagram, both host 1 and 3 are eligible after filtering. Host 1 has the highest weight and therefore has the highest priority for scheduling.
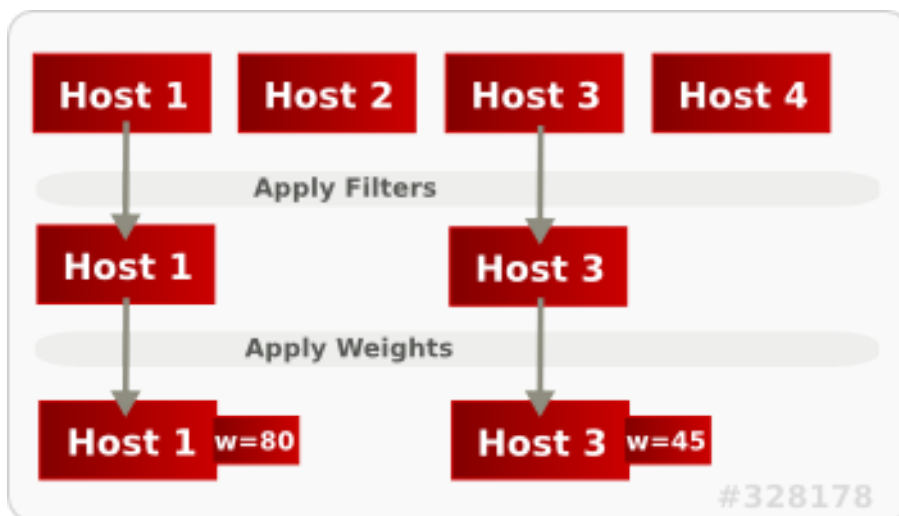
4.5.1. Configure Scheduling Filters

You define the filters you want the scheduler to use using the NovaSchedulerDefaultFilters parameter in your Compute environment file. Filters can be added or removed.

The default configuration runs the following filters in the scheduler:

- RetryFilter

- AvailabilityZoneFilter

- ComputeFilter

- ComputeCapabilitiesFilter

- ImagePropertiesFilter

- ServerGroupAntiAffinityFilter

- ServerGroupAffinityFilter

Some filters use information in parameters passed to the instance in:
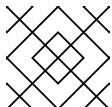
- The **nova boot** command.

- The instance's flavor (see Section 4.3.4, "Update Flavor Metadata")

- The instance's image (see Appendix A, *Image Configuration Parameters*).

All available filters are listed in the following table.

Table 4.7. Scheduling Filters

Filter	Description
AggregateImagePropert iesIsolation	Only passes hosts in host aggregates whose metadata matches the instance's image metadata; only valid if a host aggregate is specified for the instance. For more information, see Section 1.2.1, "Create an Image".
AggregateInstanceExtra SpecsFilter	Metadata in the host aggregate must match the host's flavor metadata. For more information, see Section 4.3.4, "Update Flavor Metadata".
	This filter can only be specified in the same **NovaSchedulerDefaultFilters** parameter as **ComputeCapabilitiesFilter** when you scope your flavor **extra_specs** keys by prefixing them with the correct namespace: • **ComputeCapabilitiesFilter** namespace = **"capabilities:"** • **AggregateInstanceExtraSpecsFilter** namespace = **"aggregate_instance_extra_specs:"**
AggregateMultiTenancy Isolation	A host with the specified **filter_tenant_id** can only contain instances from that tenant (project). NOTE The tenant can still place instances on other hosts.
AllHostsFilter	Passes all available hosts (however, does not disable other filters).
AvailabilityZoneFilter	Filters using the instance's specified availability zone.
ComputeCapabilitiesFilt er	Ensures Compute metadata is read correctly. Anything before the **:** is read as a namespace. For example, **quota:cpu_period** uses **quota** as the namespace and **cpu_period** as the key.
ComputeFilter	Passes only hosts that are operational and enabled.
DifferentHostFilter	Enables an instance to build on a host that is different from one or more specified hosts. Specify **different** hosts using the **nova boot** option **--different_host** option.
ImagePropertiesFilter	Only passes hosts that match the instance's image properties. For more information, see Section 1.2.1, "Create an Image".

Filter	Description
IsolatedHostsFilter	Passes only isolated hosts running isolated images that are specified using **isolated_hosts** and **isolated_images** (comma-separated values).
JsonFilter	Recognises and uses an instance's custom JSON filters: • Valid operators are: =, <, >, in, ⇐, >=, not, or, and • Recognised variables are: $free_ram_mb, $free_disk_mb, $total_usable_ram_mb, $vcpus_total, $vcpus_used
	The filter is specified as a query hint in the **nova boot** command. For example: **--hint query='['>=', '$free_disk_mb', 200 * 1024]'**
MetricFilter	Filters out hosts with unavailable metrics.
NUMATopologyFilter	Filters out hosts based on its NUMA topology. If the instance has no topology defined, any host can be used. The filter tries to match the exact NUMA topology of the instance to those of the host (it does not attempt to pack the instance onto the host). The filter also looks at the standard over-subscription limits for each NUMA node, and provides limits to the compute host accordingly.
PCIWeigher	The weigher can compute the weight based on the number of PCI devices on the host and the number of PCI devices requested by an instance. For example, if there are three hosts available, one with a single PCI device, one with multiple PCI devices and one without any PCI devices, then Compute should prioritize these hosts based on the demands of the instance. The first host should be preferred if the instance requests one PCI device, the second host if the instance requires multiple PCI devices and the third host if the instances does not request a PCI device. For more information, see Reserve NUMA Nodes with PCI Devices
RetryFilter	Filters out hosts that have failed a scheduling attempt; valid if **scheduler_max_attempts** is greater than zero (defaults to "3").
SameHostFilter	Passes one or more specified hosts; specify hosts for the instance using the **--hint same_host** option for **nova boot**.
ServerGroupAffinityFilter	Only passes hosts for a specific server group: • Give the server group the affinity policy (**nova server-group-create --policy affinity groupName**). • Build the instance with that group (**nova boot** option **--hint group=UUID**)

Filter	Description
ServerGroupAntiAffinity Filter	Only passes hosts in a server group that do not already host an instance: ● Give the server group the anti-affinity policy (**nova server-group-create --policy anti-affinity groupName**). ● Build the instance with that group (**nova boot** option **--hint group=UUID**).
SimpleCIDRAffinityFilter	Only passes hosts on the specified IP subnet range specified by the instance's cidr and **build_new_host_ip** hints. Example: **--hint build_near_host_ip=192.0.2.0 --hint cidr=/24**

4.5.2. Configure Scheduling Weights

Hosts can be weighted for scheduling; the host with the largest weight (after filtering) is selected. All weighers are given a multiplier that is applied after normalising the node's weight. A node's weight is calculated as:

w1_multiplier * norm(w1) + w2_multiplier * norm(w2) + ...

You can configure weight options in the Compute node's configuration file.

4.5.2.1. Configure Weight Options for Hosts

You can define the host weighers you would like the scheduler to use in the [DEFAULT] scheduler_weight_classes option. Valid weighers are:

● **nova.scheduler.weights.ram** - Weighs the host's available RAM.

● **nova.scheduler.weights.metrics** - Weighs the host's metrics.

● **nova.scheduler.weights.affinity** - Weighs the host's proximity to other hosts in the given server group.

● **nova.scheduler.weights.all_weighers** - Uses all host weighers (default).

Table 4.8. Host Weight Options

Weigher	Option	Description
All	[DEFAULT] **scheduler_host_subset_size**	Defines the subset size from which a host is selected (integer); must be at least **1**. A value of 1 selects the first host returned by the weighing functions. Any value less than 1 is ignored and 1 is used instead (integer value).

Weigher	Option	Description
affinity	[default] **soft_affinity_weig ht_multiplier**	Used for weighing hosts for group soft-affinity. Should be a positive floating-point number, because a negative value results in the opposite behavior, which is normally controlled by **soft_anti_affinity_weight_multiplier**.
affinity	[default] **soft_anti_affinity_ weight_multiplier**	Used for weighing hosts for group soft-anti-affinity. Should be a positive floating-point number, because a negative value results in the opposite behavior, which is normally controlled by **soft_affinity_weight_multiplier**.
metrics	[metrics] **required**	Specifies how to handle metrics in [metrics] **weight_setting** that are unavailable: • **True**- Metrics are required; if unavailable, an exception is raised. To avoid the exception, use the **MetricFilter** filter in the **scheduler_default_filters** option. • **False** - The unavailable metric is treated as a negative factor in the weighing process; the returned value is set by **weight_of_unavailable**.
metrics	[metrics] **weight_of_unavail able**	Used as the weight if any metric in [metrics] **weight_setting** is unavailable; valid if **required=False**.
metrics	[metrics] **weight_multiplier**	Mulitplier used for weighing metrics. By default, **weight_multiplier=1.0** and spreads instances across possible hosts. If this value is negative, the host with lower metrics is prioritized, and instances are stacked in hosts.

Weigher	Option	Description
metrics	[metrics] **weight_setting**	Specifies metrics and the ratio with which they are weighed; use a comma-separated list of **metric=ratio** pairs. Valid metric names are: ● **cpu.frequency** – Current CPU frequency ● **cpu.user.time** – CPU user mode time ● **cpu.kernel.time** – CPU kernel time ● **cpu.idle.time** – CPU idle time ● **cpu.iowait.time** – CPU I/O wait time ● **cpu.user.percent** – CPU user mode percentage ● **cpu.kernel.percent** – CPU kernel percentage ● **cpu.idle.percent** – CPU idle percentage ● **cpu.iowait.percent** – CPU I/O wait percentage ● **cpu.percent** - Generic CPU utilization Example: **weight_setting=cpu.user.time=1.0**
ram	[DEFAULT] **ram_weight_multiplier**	Multiplier for RAM (floating point). By default, **ram_weight_multiplier=1.0** and spreads instances across possible hosts. If this value is negative, the host with less RAM is prioritized, and instances are stacked in hosts.

4.5.3. Reserve NUMA Nodes with PCI Devices

Compute uses the filter scheduler to prioritize hosts with PCI devices for instances requesting PCI. The hosts are weighted using the **PCIWeigher** option, based on the number of PCI devices available on the host and the number of PCI devices requested by an instance. If an instance requests PCI devices, then the hosts with more PCI devices are allocated a higher weight than the others. If an instance is not requesting PCI devices, then prioritization does not take place.

This feature is especially useful in the following cases:

● As an operator, if you want to reserve nodes with PCI devices (typically expensive and with limited resources) for guest instances that request them.

● As a user launching instances, you want to ensure that PCI devices are available when required.

NOTE

For this value to be considered, one of the following values must be added to the **NovaSchedulerDefaultFilters** parameter in your Compute environment file: **PciPassthroughFilter** or **NUMATopologyFilter**.

The **pci_weight_multiplier** configuration option must be a positive value.

4.5.4. Configure Emulator Threads to run on Dedicated Physical CPU

The Compute scheduler determines the CPU resource utilization and places instances based on the number of virtual CPUs (vCPUs) in the flavor. There are a number of hypervisor operations that are performed on the host, on behalf of the guest instance, for example, with QEMU, there are threads used for the QEMU main event loop, asynchronous I/O operations and so on and these operations need to be accounted and scheduled separately.

The **libvirt** driver implements a generic placement policy for KVM which allows QEMU emulator threads to float across the same physical CPUs (pCPUs) that the vCPUs are running on. This leads to the emulator threads using time borrowed from the vCPUs operations. When you need a guest to have dedicated vCPU allocation, it is necessary to allocate one or more pCPUs for emulator threads. It is therefore necessary to describe to the scheduler any other CPU usage that might be associated with a guest and account for that during placement.

> **NOTE**
>
> In an NFV deployment, to avoid packet loss, you have to make sure that the vCPUs are never preempted.

Before you enable the emulator threads placement policy on a flavor, check that the following heat parameters are defined as follows:

- **NovaComputeCpuSharedSet**: Set this parameter to a list of CPUs defined to run emulator threads.

- **NovaSchedulerDefaultFilters**: Include **NUMATopologyFilter** in the list of defined filters.

> **NOTE**
>
> You can define or change heat parameter values on an active cluster, and then redeploy for those changes to take effect.

To isolate emulator threads, you must use a flavor configured as follows:

```
# openstack flavor set FLAVOR-NAME \
--property hw:cpu_policy=dedicated \
--property hw:emulator_threads_policy=share
```

4.6. EVACUATE INSTANCES

If you want to move an instance from a dead or shut-down compute node to a new host server in the same environment (for example, because the server needs to be swapped out), you can evacuate it using **nova evacuate**.

- An evacuation is only useful if the instance disks are on shared storage or if the instance disks are Block Storage volumes. Otherwise, the disks will not be accessible and cannot be accessed by the new compute node.

- An instance can only be evacuated from a server if the server is shut down; if the server is not shut down, the **evacuate** command will fail.

NOTE

If you have a functioning compute node, and you want to:

- Make a static copy (not running) of an instance for backup purposes or to copy the instance to a different environment, follow the procedure in Migrating VMs from an Overcloud Compute Node).

- Move an instance in a static state (not running) to a host in the same environment (shared storage not needed), migrate it using **nova migrate** (see Migrate a Static Instance).

- Move an instance in a live state (running) to a host in the same environment, migrate it using **nova live-migration** (see Migrate a Live (running) Instance).

4.6.1. Evacuate One Instance

1. Evacuate an instance using:

```
# nova evacuate [--password pass] instance_name [target_host]
```

Where:

- **--password** – Admin password to set for the evacuated instance. If a password is not specified, a random password is generated and output when evacuation is complete.

- **instance_name** – Name of the instance to be evacuated.

- **target_host** – Host to which the instance is evacuated; if you do not specify the host, the Compute scheduler selects one for you. You can find possible hosts using:

```
# nova host-list | grep compute
```

For example:

```
# nova evacuate myDemoInstance Compute2_OnEL7.myDomain
```

4.6.2. Evacuate All Instances

Evacuate all instances on a specified host using:

```
# nova host-evacuate [--target_host <target_host>] [--force] <host>
```

Where:

- **<target_host>** – The host the instance is evacuated to. If you do not specify the host, the Compute scheduler selects one for you. You can find possible hosts using the following command:

```
# nova host-list | grep compute
```

- **<host>** – Name of the host to be evacuated.

For example:

```
# nova host-evacuate --target_host Compute2_OnEL7.localdomain myDemoHost.localdomain
```

4.6.3. Configure Shared Storage

If you are using shared storage, this procedure exports the instances directory for the Compute service to the two nodes, and ensures the nodes have access. The directory path is set in the **state_path** and **instances_path** parameters in your Compute environment file. This procedure uses the default value, which is **/var/lib/nova/instances**. Only users with root access can set up shared storage.

1. **On the controller host:**

   a. Ensure the **/var/lib/nova/instances** directory has read-write access by the Compute service user (this user must be the same across controller and nodes). For example:

   ```
   drwxr-xr-x.  9 nova nova 4096 Nov  5 20:37 instances
   ```

   b. Add the following lines to the **/etc/exports** file; switch out node1_IP and node2_IP for the IP addresses of the two compute nodes:

   ```
   /var/lib/nova/instances (rw,sync,fsid=0,no_root_squash)
   /var/lib/nova/instances (rw,sync,fsid=0,no_root_squash)
   ```

   c. Export the **/var/lib/nova/instances** directory to the compute nodes.

   ```
   # exportfs -avr
   ```

   d. Restart the NFS server:

   ```
   # systemctl restart nfs-server
   ```

2. **On each compute node:**

   a. Ensure the **/var/lib/nova/instances** directory exists locally.

   b. Add the following line to the **/etc/fstab** file:

   ```
   :/var/lib/nova/instances /var/lib/nova/instances nfs4 defaults 0 0
   ```

   c. Mount the controller's instance directory (all devices listed in **/etc/fstab**):

   ```
   # mount -a -v
   ```

   d. Ensure qemu can access the directory's images:

   ```
   # ls -ld /var/lib/nova/instances
   drwxr-xr-x. 9 nova nova 4096 Nov  5 20:37 /var/lib/nova/instances
   ```

   e. Ensure that the node can see the instances directory with:

   ```
   drwxr-xr-x. 9 nova nova 4096 Nov  5 20:37 /var/lib/nova/instances
   ```

> **NOTE**
>
> You can also run the following to view all mounted devices:
>
> ```
> # df -k
> ```

4.7. MANAGE INSTANCE SNAPSHOTS

An instance snapshot allows you to create a new image from an instance. This is very convenient for upgrading base images or for taking a published image and customizing it for local use.

The difference between an image that you upload directly to the Image Service and an image that you create by snapshot is that an image created by snapshot has additional properties in the Image Service database. These properties are found in the **image_properties** table and include the following parameters:

Table 4.9. Snapshot Options

Name	Value
image_type	snapshot
instance_uuid	<uuid of instance that was snapshotted>
base_image_ref	<uuid of original image of instance that was snapshotted>
image_location	snapshot

Snapshots allow you to create new instances based on that snapshot, and potentially restore an instance to that state. Moreover, this can be performed while the instance is running.

By default, a snapshot is accessible to the users and projects that were selected while launching an instance that the snapshot is based on.

4.7.1. Create an Instance Snapshot

NOTE

If you intend to use an instance snapshot as a template to create new instances, you must ensure that the disk state is consistent. Before you create a snapshot, set the snapshot image metadata property **os_require_quiesce=yes**. For example,

```
$ glance image-update IMAGE_ID --property os_require_quiesce=yes
```

For this to work, the guest should have the **qemu-guest-agent** package installed, and the image should be created with the metadata property parameter **hw_qemu_guest_agent=yes** set. For example,

```
$ glance image-create --name NAME \
--disk-format raw \
--container-format bare \
--file FILE_NAME \
--is-public True \
--property hw_qemu_guest_agent=yes \
--progress
```

If you unconditionally enable the **hw_qemu_guest_agent=yes** parameter, then you are adding another device to the guest. This consumes a PCI slot, and will limit the number of other devices you can allocate to the guest. It also causes Windows guests to display a warning message about an unknown hardware device.

For these reasons, setting the **hw_qemu_guest_agent=yes** parameter is optional, and the parameter should be used for only those images that require the QEMU guest agent.

1. In the dashboard, select **Project > Compute > Instances**

2. Select the instance from which you want to create a snapshot.

3. In the **Actions** column, click **Create Snapshot**.

4. In the **Create Snapshot** dialog, enter a name for the snapshot and click **Create Snapshot**. The **Images** category now shows the instance snapshot.

To launch an instance from a snapshot, select the snapshot and click **Launch**.

4.7.2. Manage a Snapshot

1. In the dashboard, select **Project > Images**.

2. All snapshots you created, appear under the **Project** option.

3. For every snapshot you create, you can perform the following functions, using the dropdown list:

   a. Use the **Create Volume** option to create a volume and entering the values for volume name, description, image source, volume type, size and availability zone. For more information, see Create a Volume in the *Storage Guide*.

   b. Use the **Edit Image** option to update the snapshot image by updating the values for name, description, Kernel ID, Ramdisk ID, Architecture, Format, Minimum Disk (GB), Minimum RAM (MB), public or private. For more information, see Section 1.2.3, "Update an Image".

c. Use the **Delete Image** option to delete the snapshot.

4.7.3. Rebuild an Instance to a State in a Snapshot

In an event that you delete an instance on which a snapshot is based, the snapshot still stores the instance ID. You can check this information using the **nova image-list** command and use the snapshot to restore the instance.

1. In the dashboard, select **Project > Compute > Images**

2. Select the snapshot from which you want to restore the instance.

3. In the **Actions** column, click **Launch Instance**.

4. In the **Launch Instance** dialog, enter a name and the other details for the instance and click **Launch**.

For more information on launching an instance, see Section 4.1.2, "Launch an Instance" .

4.7.4. Consistent Snapshots

Previously, file systems had to be quiesced manually (fsfreeze) before taking a snapshot of active instances for consistent backups.

Compute's **libvirt** driver automatically requests the *QEMU Guest Agent* to freeze the file systems (and applications if **fsfreeze-hook** is installed) during an image snapshot. Support for quiescing file systems enables scheduled, automatic snapshots at the block device level.

This feature is only valid if the QEMU Guest Agent is installed (**qemu-ga**) and the image metadata enables the agent (**hw_qemu_guest_agent=yes**)

> **NOTE**
>
> Snapshots should not be considered a substitute for an actual system backup.

4.8. USE RESCUE MODE FOR INSTANCES

Compute has a method to reboot a virtual machine in rescue mode. Rescue mode provides a mechanism for access when the virtual machine image renders the instance inaccessible. A rescue virtual machine allows a user to fix their virtual machine by accessing the instance with a new root password. This feature is useful if an instance's filesystem is corrupted. By default, rescue mode starts an instance from the initial image attaching the current boot disk as a secondary one.

4.8.1. Preparing an Image for a Rescue Mode Instance

Due to the fact that both the boot disk and the disk for rescue mode have same UUID, sometimes the virtual machine can be booted from the boot disk instead of the disk for rescue mode.

To avoid this issue, you should create a new image as rescue image based on the procedure in Section 1.2.1, "Create an Image":

> **NOTE**
>
> The **rescue** image is stored in **glance** and configured in the **nova.conf** as a default, or you can select when you do the rescue.

4.8.1.1. Rescue Image if Using *ext4* Filesystem

When the base image uses **ext4** filesystem, you can create a rescue image from it using the following procedure:

1. Change the *UUID* to a random value using the **tune2fs** command:

   ```
   # tune2fs -U random /dev/DEVICE_NODE
   ```

   Here *DEVICE_NODE* is the root device node (for example, **sda**, **vda**, and so on).

2. Verify the details of the filesystem, including the new *UUID*:

   ```
   # tune2fs -l
   ```

3. Update the **/etc/fstab** to use the new *UUID*. You may need to repeat this for any additional partitions you have, that are mounted in the **fstab** by *UUID*.

4. Update the **/boot/grub2/grub.conf** file and update the *UUID* parameter with the new *UUID* of the root disk.

5. Shut down and use this image as your rescue image. This will cause the rescue image to have a new random *UUID* that will not conflict with the instance that you are rescuing.

> **NOTE**
>
> The XFS filesystem cannot change the UUID of the root device on the running virtual machine. Reboot the virtual machine until the virtual machine is launched from the disk for rescue mode.

4.8.2. Adding the Rescue Image to the OpenStack Image Service

When you have completed modifying the UUID of your image, use the following commands to add the generated rescue image to the OpenStack Image service:

1. Add the rescue image to the Image service:

   ```
   # glance image-create --name IMAGE_NAME --disk-format qcow2 \
     --container-format bare --is-public True --file IMAGE_PATH
   ```

   Here *IMAGE_NAME* is the name of the image, *IMAGE_PATH* is the location of the image.

2. Use the **image-list** command to obtain the *IMAGE_ID* required for launching an instace in the rescue mode.

   ```
   # glance image-list
   ```

You can also upload an image using the OpenStack Dashboard, see Section 1.2.2, "Upload an Image".

4.8.3. Launching an Instance in Rescue Mode

1. Since you need to rescue an instance with a specific image, rather than the default one, use the **--image** parameter:

   ```
   # nova rescue --image IMAGE_ID VIRTUAL_MACHINE_ID
   ```

   Here *IMAGE_ID* is the ID of the image you want to use and *VIRTUAL_MACHINE_ID* is ID of a virtual machine that you want to rescue.

   > **NOTE**
   >
   > The **nova rescue** command allows an instance to perform a soft shut down. This allows the guest operating system to perform a controlled shutdown before the instance is powered off. The shut down behavior is configured using **shutdown_timeout** in your Compute configuration file. The value stands for the overall period (in seconds) a guest operation system is allowed to complete the shutdown. The default timeout is 60 seconds.
   >
   > The timeout value can be overridden on a per image basis by means of **os_shutdown_timeout** that is an image metadata setting allowing different types of operating systems to specify how much time they need to shut down cleanly.

2. Reboot the virtual machine.

3. Confirm the status of the virtual machine is *RESCUE* on the controller node by using **nova list** command or by using dashboard.

4. Log in to the new virtual machine dashboard by using the password for rescue mode.

You can now make the necessary changes to your instance to fix any issues.

4.8.4. Unrescuing an Instance

You can **unrescue** the fixed instance to restart it from the boot disk.

1. Execute the following commands on the controller node.

   ```
   # nova unrescue VIRTUAL_MACHINE_ID
   ```

   Here *VIRTUAL_MACHINE_ID* is ID of a virtual machine that you want to unrescue.

The status of your instance returns to *ACTIVE* once the unrescue operation has completed successfully.

4.9. SET A CONFIGURATION DRIVE FOR INSTANCES

You can use the **config-drive** parameter to present a read-only drive to your instances. This drive can contain selected files that are then accessible to the instance. The configuration drive is attached to the instance at boot, and is presented to the instance as a partition. Configuration drives are useful when combined with *cloud-init* (for server bootstrapping), and when you want to pass large files to your instances.

4.9.1. Configuration Drive Options

Use your Compute environment file to set the following configuration drive parameters:

- **config_drive_format** – sets the format of the drive, and accepts the options **iso9660** and **vfat**. By default, it uses **iso9660**.

- **force_config_drive** – this forces the configuration drive to be presented to all instances. Set to "True".

- **mkisofs_cmd** – specifies the command to use for ISO file creation. This value must not be changed, as only *genisoimage* is supported.

4.9.2. Use a Configuration Drive

An instance attaches its configuration drive at boot time. This is enabled by the **--config-drive** option. For example, this command creates a new instance named *test-instance01* and attaches a drive containing a file named */root/user-data.txt*:

```
# nova boot --flavor m1.tiny --config-drive true --file /root/user-data.txt=/root/user-data.txt --image cirros test-instance01
```

Once the instance has booted, you can log in to it and see a file named */root/user-data.txt*.

> **NOTE**
>
> You can use the configuration drive as a source for *cloud-init* information. During the initial instance boot, *cloud-init* can automatically mount the configuration drive and run the setup scripts.

CHAPTER 5. CONFIGURING CPU PINNING WITH NUMA

This chapter describes how to use NUMA topology awareness to configure an OpenStack environment on systems with a NUMA architecture. The procedures detailed in this chapter show you how to pin virtual machines (VMs) to dedicated CPU cores, which improves scheduling and VM performance.

TIP

Background information about NUMA is available in the following article: What is NUMA and how does it work on Linux?

The following diagram provides an example of a two-node NUMA system and the way the CPU cores and memory pages are made available:
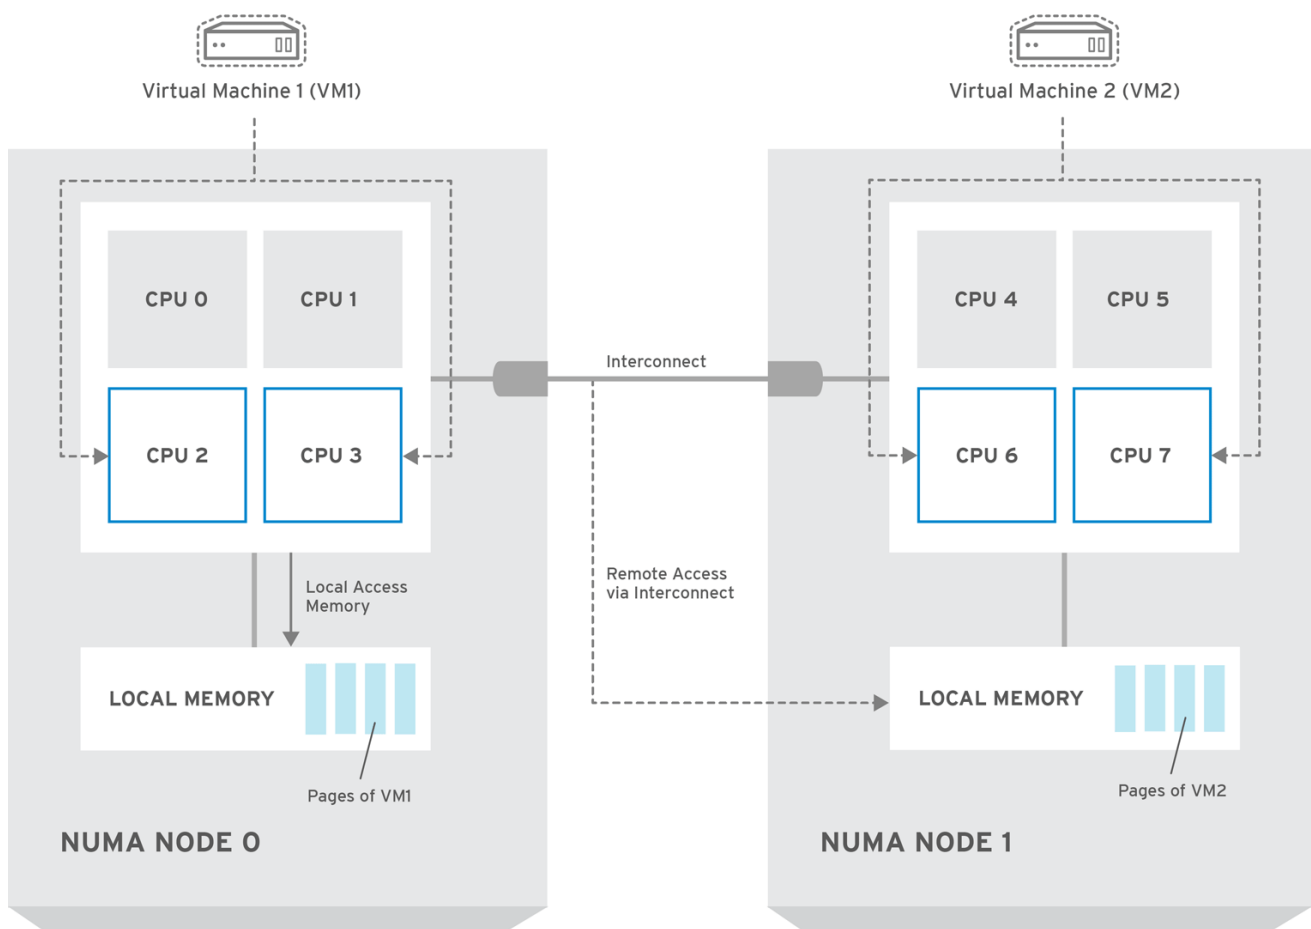
OPENSTACK_ 39825_0516

NOTE

Remote memory available via Interconnect is accessed **only** if VM1 from NUMA node 0 has a CPU core in NUMA node 1. In this case, the memory of NUMA node 1 will act as local for the third CPU core of VM1 (for example, if VM1 is allocated with CPU 4 in the diagram above), but at the same time, it will act as remote memory for the other CPU cores of the same VM.

For more details on NUMA tuning with libvirt, see the Virtualization Tuning and Optimization Guide .

5.1. COMPUTE NODE CONFIGURATION

The exact configuration depends on the NUMA topology of your host system. However, you must reserve some CPU cores across all the NUMA nodes for host processes and let the rest of the CPU cores handle your virtual machines (VMs). The following example illustrates the layout of eight CPU cores evenly spread across two NUMA nodes.

Table 5.1. Example of NUMA Topology

	Node 0		Node 1	
Host processes	Core 0	Core 1	Core 4	Core 5
VMs	Core 2	Core 3	Core 6	Core 7

NOTE

Determine the number of cores to reserve for host processes by observing the performance of the host under typical workloads.

Procedure

1. Reserve CPU cores for the VMs by setting the **NovaVcpuPinSet** configuration in the Compute environment file:

   NovaVcpuPinSet=2,3,6,7

2. Set the **NovaReservedHostMemory** option in the same file to the amount of RAM to reserve for host processes. For example, if you want to reserve 512 MB, use:

   NovaReservedHostMemory=512

3. To ensure that host processes do not run on the CPU cores reserved for VMs, set the parameter **IsolCpusList** in the Compute environment file to the CPU cores you have reserved for VMs. Specify the value of the **IsolCpusList** parameter using a list of CPU indices, or ranges separated by a whitespace. For example:

   IsolCpusList=2 3 6 7

   NOTE

   The **IsolCpusList** parameter ensures that the underlying compute node is not able to use the corresponding pCPUs for itself. The pCPUs are dedicated to the VMs.

4. To apply this configuration, deploy the overcloud:

   (undercloud) $ openstack overcloud deploy --templates \
       -e /home/stack/templates/<compute_environment_file>.yaml

5.2. SCHEDULER CONFIGURATION

Procedure

1. Open your Compute environment file.

2. Add the following values to the **NovaSchedulerDefaultFilters** parameter, if they are not already present:

   - **NUMATopologyFilter**

   - **AggregateInstanceExtraSpecsFilter**

3. Save the configuration file.

4. Deploy the overcloud.

5.3. AGGREGATE AND FLAVOR CONFIGURATION

Configure host aggregates to deploy instances that use CPU pinning on different hosts from instances that do not, to avoid unpinned instances using the resourcing requirements of pinned instances.

CAUTION

Do not deploy instances with NUMA topology on the same hosts as instances that do not have NUMA topology.

Prepare your OpenStack environment for running virtual machine instances pinned to specific resources by completing the following steps on a system with the Compute CLI.

Procedure

1. Load the **admin** credentials:

   ```
   source ~/keystonerc_admin
   ```

2. Create an aggregate for the hosts that will receive pinning requests:

   ```
   nova aggregate-create <aggregate-name-pinned>
   ```

3. Enable the pinning by editing the metadata for the aggregate:

   ```
   nova aggregate-set-metadata <aggregate-pinned-UUID> pinned=true
   ```

4. Create an aggregate for other hosts:

   ```
   nova aggregate-create <aggregate-name-unpinned>
   ```

5. Edit the metadata for this aggregate accordingly:

   ```
   nova aggregate-set-metadata <aggregate-unpinned-UUID> pinned=false
   ```

6. Change your existing flavors' specifications to this one:

```
for i in $(nova flavor-list | cut -f 2 -d ' ' | grep -o '[0-9]*'); do nova flavor-key $i set
"aggregate_instance_extra_specs:pinned"="false"; done
```

7. Create a flavor for the hosts that will receive pinning requests:

```
nova flavor-create <flavor-name-pinned> <flavor-ID> <RAM> <disk-size> <vCPUs>
```

Where:

- **<flavor-ID>** – Set to **auto** if you want **nova** to generate a UUID.

- **<RAM>** – Specify the required RAM in MB.

- **<disk-size>** – Specify the required disk size in GB.

- **<vCPUs>** – The number of virtual CPUs that you want to reserve.

8. Set the **hw:cpu_policy** specification of this flavor to **dedicated** so as to require dedicated resources, which enables CPU pinning, and also the **hw:cpu_thread_policy** specification to **require**, which places each vCPU on thread siblings:

```
nova flavor-key <flavor-name-pinned> set hw:cpu_policy=dedicated
nova flavor-key <flavor-name-pinned> set hw:cpu_thread_policy=require
```

> **NOTE**
>
> If the host does not have an SMT architecture or enough CPU cores with free thread siblings, scheduling will fail. If such behavior is undesired, or if your hosts simply do not have an SMT architecture, do not use the **hw:cpu_thread_policy** specification, or set it to **prefer** instead of **require**. The (default) **prefer** policy ensures that thread siblings are used when available.

9. Set the **aggregate_instance_extra_specs:pinned** specification to "true" to ensure that instances based on this flavor have this specification in their aggregate metadata:

```
nova flavor-key <flavor-name-pinned> set aggregate_instance_extra_specs:pinned=true
```

10. Add some hosts to the new aggregates:

```
nova aggregate-add-host <aggregate-pinned-UUID> <host_name>
nova aggregate-add-host <aggregate-unpinned-UUID> <host_name>
```

11. Boot an instance using the new flavor:

```
nova boot --image <image-name> --flavor <flavor-name-pinned> <server-name>
```

12. To verify that the new server has been placed correctly, run the following command and check for **OS-EXT-SRV-ATTR:hypervisor_hostname** in the output:

```
nova show <server-name>
```

CHAPTER 6. CONFIGURING VIRTUAL GPU FOR GUEST INSTANCES

> **IMPORTANT**
>
> This feature is available in this release as a *Technology Preview*, and therefore is not fully supported by Red Hat. It should only be used for testing, and should not be deployed in a production environment. For more information about Technology Preview features, see Scope of Coverage Details.

To support GPU-based rendering on your guest instances, you can define and manage virtual GPU (vGPU) resources according to your available physical GPU devices and your hypervisor type. This configuration allows you to divide the rendering workloads between all your physical GPU devices more effectively, and to have more control over scheduling, tuning, and monitoring your vGPU-enabled guest instances.

To enable vGPU in OpenStack Compute, you create flavors that you can use to request Red Hat Enterprise Linux guests with vGPU devices, and then you assign those flavors to Compute instances. Each instance can then support GPU workloads with virtual GPU devices that correspond to the physical GPU devices.

The OpenStack Compute service tracks the number and size of the vGPU devices that are available on each host, schedules guests to these hosts based on the flavor, attaches the devices, and monitors usage on an ongoing basis. In case the guest is no longer available, OpenStack Compute adds the vGPU devices back to the available pool.

6.1. SUPPORTED CONFIGURATIONS AND LIMITATIONS

This section lists currently supported virtual GPU (vGPU) graphics cards, as well as considerations and limitations for setting up vGPU devices in OpenStack Compute.

Supported GPU cards

The following table lists the currently supported GPU cards for Red Hat Enterprise Linux and KVM.

Vendor	Model	License
NVIDIA	Tesla M6	NVIDIA vGPU GRID
NVIDIA	Tesla M10	NVIDIA vGPU GRID
NVIDIA	Tesla M60	NVIDIA vGPU GRID
NVIDIA	Tesla P4	NVIDIA vGPU GRID
NVIDIA	Tesla P6	NVIDIA GRID
NVIDIA	Tesla P40	NVIDIA vGPU GRID
NVIDIA	Tesla P100 PCIe 16 GB	NVIDIA vGPU GRID

Vendor	Model	License
NVIDIA	Tesla P100 SXM2 16 GB	NVIDIA vGPU GRID
NVIDIA	Tesla PCIe 12GB	NVIDIA vGPU GRID
NVIDIA	Tesla V100 SXM2	NVIDIA vGPU GRID
NVIDIA	Tesla V100 PCIe	NVIDIA vGPU GRID
NVIDIA	Tesla V100 FHHL	NVIDIA vGPU GRID

Supported vGPU types

The following table lists currently supported vGPU types and their typical specifications and license types.

Vendor	Virtual GPU Type	Frame Buffer (GB)	Maximum Resolution Per Display Head	Maximum vGPUs per pGPU	Maximum vGPUs per board (M10 = 4 x pGPU)	License
NVIDIA	M10-0B	0.512	2560x1600	16	64	GRID Virtual PC or Quadro vDWS
NVIDIA	M10-2B	2	4096x2160	4	16	GRID Virtual PC or Quadro vDWS
NVIDIA	M10-1B	1	2560x1600	8	32	GRID Virtual PC or Quadro vDWS
NVIDIA	M10-8A	8	1280x1024	1	4	GRID Virtual PC or Quadro vDWS

Limitations and considerations

- You can use only one vGPU type for each Compute host.

- You can use only one vGPU resource for each Compute instance.

- Live migration of vGPU between hosts is not supported.

- Suspend operations on a vGPU-enabled guest is not supported due to a libvirt limitation. Instead, you can snapshot or shelve the instance.

- Resize and cold migration operations on an instance with a vGPU flavor does not automatically re-allocate the vGPU resources to the instance. After you resize or migrate the instance, you must rebuild it manually to re-allocate the vGPU resources.

- By default, vGPU types on Compute hosts are not exposed to API users. To allow access, you can add the hosts to a host aggregate. For general information about host aggregates, see Section 4.4, "Manage Host Aggregates"

- If you use NVIDIA accelerator hardware, you must comply with the NVIDIA licensing requirements. For example, NVIDIA vGPU GRID requires a licensing server. For more information about the NVIDIA licensing requirements, see the NVIDIA License Server Release Notes web page.

6.2. DEPLOYING NVIDIA GRID VGPU

This section describes how to deploy virtual GPU (vGPU) for NVIDIA devices on your Compute node hosts and on your guest instances. This end-to-end process includes the following steps:

1. Building a custom GPU-enabled overcloud image

2. Preparing the GPU role, profile, and flavor

3. Configuring and deploying the overcloud

4. Building a custom vGPU-enabled guest image

5. Preparing the vGPU flavor for the instances

6. Launching and configuring the vGPU-enabled instances

Prerequisites

Before you deploy NVIDIA GRID vGPU on your overcloud, make sure that your environment meets the following requirements:

- Your deployment must meet the requirements for vGPU devices, as described in Section 6.1, "Supported configurations and limitations".

- Your undercloud must be deployed and the default overcloud image must be uploaded to Glance.

- You must comply with the NVIDIA GRID licensing requirements and you must have the URL of your self-hosted license server. For more information about the NVIDIA licensing requirements and self-hosted server installation, see the NVIDIA License Server Release Notes web page.

6.2.1. Build a custom GPU overcloud image

Perform the following steps on the undercloud to install the NVIDIA GRID host driver on an overcloud Compute image and upload the image to Glance.

1. Copy the overcloud image and add the **gpu** suffix to the copied image.

   ```
   $ cp overcloud-full.qcow2 overcloud-full-gpu.qcow2
   ```

2. Install an ISO image generator tool from YUM.

```
$ sudo yum install genisoimage -y
```

3. Download the NVIDIA GRID host driver RPM package that corresponds to your GPU device from the NVIDIA website. To determine which driver you need, see the NVIDIA Driver Downloads Portal.

> **NOTE**
>
> You must be a registered NVIDIA customer to download the drivers from the portal.

4. Create an ISO image from the driver RPM package and save the image in the *nvidia-guest* directory. You will use this ISO image to install the driver on your Compute nodes in subsequent steps.

```
$ genisoimage -o nvidia-guest.iso -R -J -V NVIDIA nvidia-guest/
I: -input-charset not specified, using utf-8 (detected in locale settings)
  9.06% done, estimate finish Wed Oct 31 11:24:46 2018
 18.08% done, estimate finish Wed Oct 31 11:24:46 2018
 27.14% done, estimate finish Wed Oct 31 11:24:46 2018
 36.17% done, estimate finish Wed Oct 31 11:24:46 2018
 45.22% done, estimate finish Wed Oct 31 11:24:46 2018
 54.25% done, estimate finish Wed Oct 31 11:24:46 2018
 63.31% done, estimate finish Wed Oct 31 11:24:46 2018
 72.34% done, estimate finish Wed Oct 31 11:24:46 2018
 81.39% done, estimate finish Wed Oct 31 11:24:46 2018
 90.42% done, estimate finish Wed Oct 31 11:24:46 2018
 99.48% done, estimate finish Wed Oct 31 11:24:46 2018
Total translation table size: 0
Total rockridge attributes bytes: 358
Total directory bytes: 0
Path table size(bytes): 10
Max brk space used 0
55297 extents written (108 MB)
```

5. Create a driver installation script for your Compute nodes. This script installs the NVIDIA GRID host driver on each Compute node that you run it on. In this example the script is named *install_nvidia.sh*.

```
#/bin/bash

# NVIDIA GRID package
mkdir /tmp/mount
mount LABEL=NVIDIA /tmp/mount
rpm -ivh /tmp/mount/NVIDIA-vGPU-rhel-7.6-410.91.x86_64.rpm
```

6. Customize the overcloud image by attaching the ISO image that you generated and running the driver installation script that you created.

```
$ virt-customize --attach nvidia-packages.iso -a overcloud-full-gpu.qcow2  -v --run
install_nvidia.sh
[   0.0] Examining the guest ...
```

```
libguestfs: launch: program=virt-customize
libguestfs: launch: version=1.36.10rhel=7,release=6.el7_5.2,libvirt
libguestfs: launch: backend registered: unix
libguestfs: launch: backend registered: uml
libguestfs: launch: backend registered: libvirt
```

7. Relabel the customized image with SELinux.

```
$ virt-customize -a overcloud-full-gpu.qcow2 --selinux-relabel
[   0.0] Examining the guest ...
[   2.2] Setting a random seed
[   2.2] SELinux relabelling
[  27.4] Finishing off
```

8. Prepare the custom image files for a Glance upload.

```
$ mkdir /var/image/x86_64/image
$ guestmount -a overcloud-full-gpu.qcow2 -i --ro image
$ cp image/boot/vmlinuz-3.10.0-862.14.4.el7.x86_64 ./overcloud-full-gpu.vmlinuz
$ cp image/boot/initramfs-3.10.0-862.14.4.el7.x86_64.img ./overcloud-full-gpu.initrd
```

9. From the undercloud, upload the custom image to Glance.

```
(undercloud) $ openstack overcloud image upload --update-existing --os-image-name
overcloud-full-gpu.qcow2
```

6.2.2. Configure the vGPU role, profile, and flavor

After you build the custom GPU overcloud image, you prepare the Compute nodes for GPU-enabled overcloud deployment. This section describes how to configure the role, profile, and flavor for the GPU-enabled Compute nodes.

1. Create the new **ComputeGPU** role file by copying the file
/home/stack/templates/roles/Compute.yaml to
/home/stack/templates/roles/ComputeGPU.yaml and editing the following file sections:

Table 6.1. ComputeGPU role file edits

Section	Current value	New value
Role comment	Role: Compute	Role: ComputeGpu
Role name	name: Compute	name: ComputeGpu
Description	Basic Compute Node role	GPU role
CountDefault	1	0
ImageDefault	overcloud-full	overcloud-gpu
HostnameFormatDefault	-compute-	-computegpu-

Section	Current value	New value
deprecated_nic_config_name	compute.yaml	compute-gpu.yaml

2. Generate a new roles data file named *gpu_roles_data.yaml* that includes the **Controller**, **Compute**, and **ComputeGpu** roles.

    (undercloud) [stack@director templates]$ openstack overcloud roles generate -o /home/stack/templates/gpu_roles_data.yaml Controller Compute ComputeGpu

    The following example shows the **ComputeGpu** role details:

    ```
    ###############################################################################
    # Role: ComputeGpu                                     #
    ###############################################################################
    - name: ComputeGpu
      description: |
        GPU Compute Node role
      CountDefault: 1
      ImageDefault: overcloud-gpu
      networks:
        - InternalApi
        - Tenant
        - Storage
      HostnameFormatDefault: '%stackname%-computegpu-%index%'
      RoleParametersDefault:
        TunedProfileName: "virtual-host"
      # Deprecated & backward-compatible values (FIXME: Make parameters consistent)
      # Set uses_deprecated_params to True if any deprecated params are used.
      uses_deprecated_params: True
      deprecated_param_image: 'NovaImage'
      deprecated_param_extraconfig: 'NovaComputeExtraConfig'
      deprecated_param_metadata: 'NovaComputeServerMetadata'
      deprecated_param_scheduler_hints: 'NovaComputeSchedulerHints'
      deprecated_param_ips: 'NovaComputeIPs'
      deprecated_server_resource_name: 'NovaCompute'
      deprecated_nic_config_name: 'compute-gpu.yaml'
      ServicesDefault:
        - OS::TripleO::Services::Aide
        - OS::TripleO::Services::AuditD
        - OS::TripleO::Services::CACerts
        - OS::TripleO::Services::CephClient
        - OS::TripleO::Services::CephExternal
        - OS::TripleO::Services::CertmongerUser
        - OS::TripleO::Services::Collectd
        - OS::TripleO::Services::ComputeCeilometerAgent
        - OS::TripleO::Services::ComputeNeutronCorePlugin
        - OS::TripleO::Services::ComputeNeutronL3Agent
        - OS::TripleO::Services::ComputeNeutronMetadataAgent
        - OS::TripleO::Services::ComputeNeutronOvsAgent
        - OS::TripleO::Services::Docker
        - OS::TripleO::Services::Fluentd
        - OS::TripleO::Services::Ipsec
        - OS::TripleO::Services::Iscsid
    ```

```
    - OS::TripleO::Services::Kernel
    - OS::TripleO::Services::LoginDefs
    - OS::TripleO::Services::MetricsQdr
    - OS::TripleO::Services::MySQLClient
    - OS::TripleO::Services::NeutronBgpVpnBagpipe
    - OS::TripleO::Services::NeutronLinuxbridgeAgent
    - OS::TripleO::Services::NeutronVppAgent
    - OS::TripleO::Services::NovaCompute
    - OS::TripleO::Services::NovaLibvirt
    - OS::TripleO::Services::NovaLibvirtGuests
    - OS::TripleO::Services::NovaMigrationTarget
    - OS::TripleO::Services::Ntp
    - OS::TripleO::Services::ContainersLogrotateCrond
    - OS::TripleO::Services::OpenDaylightOvs
    - OS::TripleO::Services::Rhsm
    - OS::TripleO::Services::RsyslogSidecar
    - OS::TripleO::Services::Securetty
    - OS::TripleO::Services::SensuClient
    - OS::TripleO::Services::SkydiveAgent
    - OS::TripleO::Services::Snmp
    - OS::TripleO::Services::Sshd
    - OS::TripleO::Services::Timezone
    - OS::TripleO::Services::TripleoFirewall
    - OS::TripleO::Services::TripleoPackages
    - OS::TripleO::Services::Tuned
    - OS::TripleO::Services::Vpp
    - OS::TripleO::Services::OVNController
    - OS::TripleO::Services::OVNMetadataAgent
    - OS::TripleO::Services::Ptp
```

3. Create the **compute-vgpu-nvidia** flavor to tag nodes that you want to designate for vGPU workloads.

```
(undercloud) [stack@director templates]$ openstack flavor create --id auto --ram 6144 --disk
40 --vcpus 4 compute-vgpu-nvidia
+----------------------------+--------------------------------------+
| Field                      | Value                                |
+----------------------------+--------------------------------------+
| OS-FLV-DISABLED:disabled   | False                                |
| OS-FLV-EXT-DATA:ephemeral  | 0                                    |
| disk                       | 40                                   |
| id                         | 9cb47954-be00-47c6-a57f-44db35be3e69 |
| name                       | compute-vgpu-nvidia                  |
| os-flavor-access:is_public | True                                 |
| properties                 |                                      |
| ram                        | 6144                                 |
| rxtx_factor                | 1.0                                  |
| swap                       |                                      |
| vcpus                      | 4                                    |
+----------------------------+--------------------------------------+
```

4. Tag each node that you want to designate for GPU workloads with the **compute-vgpu-nvidia** profile.

```
(undercloud) [stack@director templates]$ openstack baremetal node set --property
capabilities='profile:compute-vgpu-nvidia,boot_option:local' 9d07a673-b6bf-4a20-a538-
3b05e8fa2c13
```

5. Register the overcloud and run the standard hardware introspection on your nodes.

6.2.3. Prepare configuration files and deploying the overcloud

After you prepare your overcloud for vGPU, you retrieve and assign the vGPU type that corresponds to the physical GPU device in your environment and prepare the configuration templates.

Configure the vGPU type for your NVIDIA device

To determine the vGPU type for your physical GPU device, you must check the available device type from a different machine. You can perform these steps from any temporary Red Hat Enterprise Linux unused Compute node, and then delete the node. You do not need to deploy the overcloud to perform these steps.

1. Install Red Hat Enterprise Linux and the NVIDIA GRID driver on one Compute node and launch the node. For information on installing the NVIDIA GRID driver, see Section 6.2.1, "Build a custom GPU overcloud image".

2. On the Compute node, locate the vGPU type of the physical GPU device that you want to enable. For libvirt, virtual GPUs are seen as mediated devices, or **mdev** type devices. To discover the supported **mdev** devices, run the following command:

```
[root@overcloud-computegpu-0 ~]# ls
/sys/class/mdev_bus/0000\:06\:00.0/mdev_supported_types/
nvidia-11 nvidia-12 nvidia-13 nvidia-14 nvidia-15 nvidia-16 nvidia-17 nvidia-18 nvidia-19
nvidia-20 nvidia-21 nvidia-210 nvidia-22

[root@overcloud-computegpu-0 ~]# cat
/sys/class/mdev_bus/0000\:06\:00.0/mdev_supported_types/nvidia-18/description
num_heads=4, frl_config=60, framebuffer=2048M, max_resolution=4096x2160,
max_instance=4
```

Prepare the configuration templates

1. Add the *compute-gpu.yaml* file to the *network-environment.yaml* file. For example:

```
resource_registry:
  OS::TripleO::Compute::Net::SoftwareConfig: /home/stack/templates/nic-
configs/compute.yaml
  OS::TripleO::ComputeGpu::Net::SoftwareConfig: /home/stack/templates/nic-
configs/compute-gpu.yaml
  OS::TripleO::Controller::Net::SoftwareConfig: /home/stack/templates/nic-
configs/controller.yaml
  #OS::TripleO::AllNodes::Validation: OS::Heat::None
```

2. Add the **OvercloudComputeGpuFlavor** flavor to the *node-info.yaml* file. For example:

```
parameter_defaults:
  OvercloudControllerFlavor: control
  OvercloudComputeFlavor: compute
```

```
OvercloudComputeGpuFlavor: compute-vgpu-nvidia
ControllerCount: 1
ComputeCount: 0
ComputeGpuCount: 1
NtpServer: `NTP_SERVER_URL`
NeutronNetworkType: vxlan,vlan
NeutronTunnelTypes: vxlan
```

Replace the **NTP_SERVER_URL** variable with the address of your NTP server.

3. Create a *gpu.yaml* file with the vGPU type that you retrieved for your GPU device. For example:

```
parameter_defaults:
  ComputeGpuExtraConfig:
    nova::compute::vgpu::enabled_vgpu_types:
      - nvidia-18
```

> **NOTE**
>
> Only one virtual GPU type is supported per physical GPU. If you specify multiple vGPU types in this property, only the first type is used.

Deploy the overcloud

Run the **overcloud deploy** command with the custom GPU image and the configuration templates that you prepared.

```
$ openstack overcloud deploy -r /home/stack/templates/nvidia/gpu_roles_data.yaml  -e
/home/stack/templates/nvidia/gpu.yaml
```

6.2.4. Build a custom GPU guest image

After you deploy the overcloud with GPU-enabled Compute nodes, you build a custom vGPU-enabled instance image with the NVIDIA GRID guest driver and license file.

Create the NVIDIA GRID license file

In the overcloud host, create a *gridd.conf* file that contains the NVIDIA GRID license information. Use the license server information from your self-hosted NVIDIA GRID license server that you installed previously. For example:

```
# /etc/nvidia/gridd.conf.template - Configuration file for NVIDIA Grid Daemon

# This is a template for the configuration file for NVIDIA Grid Daemon.
# For details on the file format, please refer to the nvidia-gridd(1)
# man page.

# Description: Set License Server Address
# Data type: string
# Format:  "<address>"
ServerAddress=[NVIDIA_LICENSE_SERVER_URL]

# Description: Set License Server port number
# Data type: integer
```

```
# Format:  <port>, default is 7070
ServerPort=[PORT_NUMBER]

# Description: Set Backup License Server Address
# Data type: string
# Format:  "<address>"
#BackupServerAddress=

# Description: Set Backup License Server port number
# Data type: integer
# Format:  <port>, default is 7070
#BackupServerPort=

# Description: Set Feature to be enabled
# Data type: integer
# Possible values:
#    0 => for unlicensed state
#    1 => for GRID vGPU
#    2 => for Quadro Virtual Datacenter Workstation
FeatureType=[TYPE_ID]

# Description: Parameter to enable or disable Grid Licensing tab in nvidia-settings
# Data type: boolean
# Possible values: TRUE or FALSE, default is FALSE
EnableUI=TRUE

# Description: Set license borrow period in minutes
# Data type: integer
# Possible values: 10 to 10080 mins(7 days), default is 1440 mins(1 day)
#LicenseInterval=1440

# Description: Set license linger period in minutes
# Data type: integer
# Possible values: 0 to 10080 mins(7 days), default is 0 mins
#LingerInterval=10
```

Prepare the guest image and the NVIDIA GRID guest driver

1. Download the NVIDIA GRID guest driver RPM package that corresponds to your GPU device from the NVIDIA website. To determine which driver you need, see the NVIDIA Driver Downloads Portal.

   NOTE

   You must be a registered NVIDIA customer to download the drivers from the portal.

2. Create an ISO image from the driver RPM package. You will use this ISO image to install the driver on your Compute nodes in subsequent steps.

   ```
   [root@virtlab607 guest]# genisoimage -o nvidia-guest.iso -R -J -V NVIDIA nvidia-guest/
   I: -input-charset not specified, using utf-8 (detected in locale settings)
     9.06% done, estimate finish Wed Oct 31 10:59:50 2018
    18.08% done, estimate finish Wed Oct 31 10:59:50 2018
    27.14% done, estimate finish Wed Oct 31 10:59:50 2018
   ```

```
36.17% done, estimate finish Wed Oct 31 10:59:50 2018
45.22% done, estimate finish Wed Oct 31 10:59:50 2018
54.25% done, estimate finish Wed Oct 31 10:59:50 2018
63.31% done, estimate finish Wed Oct 31 10:59:50 2018
72.34% done, estimate finish Wed Oct 31 10:59:50 2018
81.39% done, estimate finish Wed Oct 31 10:59:50 2018
90.42% done, estimate finish Wed Oct 31 10:59:50 2018
99.48% done, estimate finish Wed Oct 31 10:59:50 2018
Total translation table size: 0
Total rockridge attributes bytes: 358
Total directory bytes: 0
Path table size(bytes): 10
Max brk space used 0
55297 extents written (108 MB)
```

3. Copy the guest image that you want to customize for GPU instances. For example:

```
[root@virtlab607 guest]# cp rhel-server-7.6-update-4-x86_64-kvm.qcow2 rhel-server-7.6-update-4-x86_64-kvm-gpu.qcow2
```

Create and run the customization script

By default, you must install the NVIDIA GRID drivers on each instance that you want to designate for GPU workloads. This process involves modifying the guest image, rebooting, and then installing the guest drivers. You can create a script to automate this process for the guest instances.

1. Create a script named *nvidia-prepare-guest.sh* to enable the required repositories, update the instance to the latest kernel, install the NVIDIA GRID guest driver, and attach the *gridd.conf* license file to the instance.

```
#/bin/bash

# Add build tooling
subscription-manager register --username [USERNAME] --password [PASSWORD]
subscription-manager attach --pool=8a85f98c651a88990165399d8eea03e7
subscription-manager repos --disable=*
subscription-manager repos --enable=rhel-7-server-rpms
yum upgrade -y
yum install -y gcc make kernel-devel cpp glibc-devel glibc-headers kernel-headers libmpc mpfr

# NVIDIA GRID guest script
mkdir /tmp/mount
mount LABEL=NVIDIA /tmp/mount
/bin/sh /tmp/mount/NVIDIA-Linux-x86_64-410.71-grid.run

mkdir -p /etc/nvidia
cp /tmp/mount/gridd.conf /etc/nvidia
```

2. Run the script on the guest image that you copied previously. For example:

```
$ virt-customize --attach nvidia-guest.iso -a rhel-server-7.6-update-4-x86_64-kvm-gpu.qcow2 -v --run nvidia-prepare-guest.sh
```

3. Upload the custom guest image to Glance.

```
(overcloud) [stack@director ~]$ openstack image create rhelgpu --file
/var/images/x86_64/rhel-server-7.6-update-4-x86_64-kvm-gpu.qcow2 --disk-format qcow2 --
container-format bare --public
```

6.2.5. Create a vGPU profile for instances

After you build the custom guest image, you create a GPU flavor and assign a vGPU resource to that flavor. When you later launch instances with this flavor, the vGPU resource will be available to each instance.

> **NOTE**
>
> You can assign only one vGPU resource for each instance.

1. Create an NVIDIA GPU flavor to tag each instance that you want to designate for GPU workloads. For example:

```
(overcloud) [stack@virtlab-director2 ~]$ openstack flavor create --vcpus 6 --ram 8192 --disk
100 m1.small-gpu
+---------------------------+-------------------------------------+
| Field                     | Value                               |
+---------------------------+-------------------------------------+
| OS-FLV-DISABLED:disabled  | False                               |
| OS-FLV-EXT-DATA:ephemeral | 0                                   |
| disk                      | 100                                 |
| id                        | a27b14dd-c42d-4084-9b6a-225555876f68 |
| name                      | m1.small-gpu                        |
| os-flavor-access:is_public | True                               |
| properties                |                                     |
| ram                       | 8192                                |
| rxtx_factor               | 1.0                                 |
| swap                      |                                     |
| vcpus                     | 6                                   |
+---------------------------+-------------------------------------+
```

2. Assign a vGPU resource to the flavor that you created. Currently you can assign only one vGPU for each instance.

```
(overcloud) [stack@virtlab-director2 ~]$ openstack flavor set m1.small-gpu --property
"resources:VGPU=1"

(overcloud) [stack@virtlab-director2 ~]$ openstack flavor show m1.small-gpu
+---------------------------+-------------------------------------+
| Field                     | Value                               |
+---------------------------+-------------------------------------+
| OS-FLV-DISABLED:disabled  | False                               |
| OS-FLV-EXT-DATA:ephemeral | 0                                   |
| access_project_ids        | None                                |
| disk                      | 100                                 |
| id                        | a27b14dd-c42d-4084-9b6a-225555876f68 |
| name                      | m1.small-gpu                        |
| os-flavor-access:is_public | True                               |
| properties                | resources:VGPU='1'                  |
| ram                       | 8192                                |
```

```
| rxtx_factor            | 1.0                    |
| swap                   |                        |
| vcpus                  | 6                      |
+------------------------+------------------------+
```

6.2.6. Launch and test a vGPU instance

After you prepare the guest image and create the GPU flavor, you launch the GPU-enabled instance and install the NVIDIA guest driver from the ISO that you attached to the custom image in Section 6.2.4, "Build a custom GPU guest image".

1. Launch a new instance with the GPU flavor that you created in Section 6.2.5, "Create a vGPU profile for instances". For example:

   ```
   (overcloud) [stack@virtlab-director2 ~]$ openstack server create --flavor m1.small-gpu --image rhelgpu --security-group web --nic net-id=internal0 --key-name lambda instance0
   ```

2. Log in to the instance and install the NVIDIA GRID driver. The exact installer name is available from the files that you attached to the guest image. For example:

   ```
   [root@instance0 tmp]# sh NVIDIA-Linux-x86_64-410.71-grid.run
   ```

3. Check the status of the NVIDIA GRID daemon.

   ```
   [root@instance0 nvidia]# systemctl status nvidia-gridd.service
   ● nvidia-gridd.service - NVIDIA Grid Daemon
      Loaded: loaded (/usr/lib/systemd/system/nvidia-gridd.service; enabled; vendor preset: disabled)
      Active: active (running) since Wed 2018-10-31 20:00:41 EDT; 15s ago
    Process: 18143 ExecStopPost=/bin/rm -rf /var/run/nvidia-gridd (code=exited, status=0/SUCCESS)
    Process: 18145 ExecStart=/usr/bin/nvidia-gridd (code=exited, status=0/SUCCESS)
    Main PID: 18146 (nvidia-gridd)
      CGroup: /system.slice/nvidia-gridd.service
              └─18146 /usr/bin/nvidia-gridd

   Oct 31 20:00:41 instance0 systemd[1]: Stopped NVIDIA Grid Daemon.
   Oct 31 20:00:41 instance0 systemd[1]: Starting NVIDIA Grid Daemon...
   Oct 31 20:00:41 instance0 systemd[1]: Started NVIDIA Grid Daemon.
   Oct 31 20:00:41 instance0 nvidia-gridd[18146]: Started (18146)
   Oct 31 20:00:41 instance0 nvidia-gridd[18146]: Ignore Service Provider Licensing.
   Oct 31 20:00:41 instance0 nvidia-gridd[18146]: Calling load_byte_array(tra)
   Oct 31 20:00:42 instance0 nvidia-gridd[18146]: Acquiring license for GRID vGPU Edition.
   Oct 31 20:00:42 instance0 nvidia-gridd[18146]: Calling load_byte_array(tra)
   Oct 31 20:00:45 instance0 nvidia-gridd[18146]: License acquired successfully. (Info: http://dhcp158-15.virt.lab.eng.bos.redhat.com:7070/request; GRID-Virtual-WS,2.0)
   ```

APPENDIX A. IMAGE CONFIGURATION PARAMETERS

The following keys can be used with the **property** option for both the **glance image-update** and **glance image-create** commands.

```
$ glance image-update IMG-UUID --property architecture=x86_64
```

NOTE

Behavior set using image properties overrides behavior set using flavors. For more information, see Section 4.3, "Manage Flavors".

Table A.1. Property Keys

Specific to	Key	Description	Supported values
All	architecture	The CPU architecture that must be supported by the hypervisor. For example, **x86_64**, **arm**, or **ppc64**. Run **uname -m** to get the architecture of a machine. We strongly recommend using the architecture data vocabulary defined by the libosinfo project for this purpose.	- alpha-DEC 64-bit RISC - armv7l-ARM Cortex-A7 MPCore - cris-Ethernet, Token Ring, AXis-Code Reduced Instruction Set - i686-Intel sixth-generation x86 (P6 micro architecture) - ia64-Itanium - lm32-Lattice Micro32 - m68k-Motorola 68000 - microblaze-Xilinx 32-bit FPGA (Big Endian) - microblazeel-Xilinx 32-bit FPGA (Little Endian) - mips-MIPS 32-bit RISC (Big Endian) - mipsel-MIPS 32-bit RISC (Little Endian) - mips64-MIPS 64-bit RISC (Big Endian) - mips64el-MIPS 64-bit RISC (Little Endian) - openrisc-OpenCores RISC - parisc-HP Precision Architecture RISC - parisc64-HP Precision Architecture 64-bit RISC - ppc-PowerPC 32-bit - ppc64-PowerPC 64-bit

Specific to	Key	Description	Supported values
			• ppcemb-PowerPC (Embedded 32-bit) • s390-IBM Enterprise Systems Architecture/390 • s390x-S/390 64-bit • sh4-SuperH SH-4 (Little Endian) • sh4eb-SuperH SH-4 (Big Endian) • sparc-Scalable Processor Architecture, 32-bit • sparc64-Scalable Processor Architecture, 64-bit • unicore32-Microprocessor Research and Development Center RISC Unicore32 • x86_64-64-bit extension of IA-32 • xtensa-Tensilica Xtensa configurable microprocessor core • xtensaeb-Tensilica Xtensa configurable microprocessor core (Big Endian)
All	hypervisor_type	The hypervisor type.	**kvm**, **vmware**
All	instance_uuid	For snapshot images, this is the UUID of the server used to create this image.	Valid server UUID
All	kernel_id	The ID of an image stored in the Image Service that should be used as the kernel when booting an AMI-style image.	Valid image ID

Specific to	Key	Description	Supported values
All	os_distro	The common name of the operating system distribution in lowercase (uses the same data vocabulary as the libosinfo project). Specify only a recognized value for this field. Deprecated values are listed to assist you in searching for the recognized value.	- arch-Arch Linux. Do not use archlinux or org.archlinux - centos-Community Enterprise Operating System. Do not use org.centos or CentOS - debian-Debian. Do not use Debian or org.debian - fedora-Fedora. Do not use Fedora, org.fedora, or org.fedoraproject - freebsd-FreeBSD. Do not use org.freebsd, freeBSD, or FreeBSD - gentoo-Gentoo Linux. Do not use Gentoo or org.gentoo - mandrake-Mandrakelinux (MandrakeSoft) distribution. Do not use mandrakelinux or MandrakeLinux - mandriva-Mandriva Linux. Do not use mandrivalinux - mes-Mandriva Enterprise Server. Do not use mandrivaent or mandrivaES - msdos-Microsoft Disc Operating System. Do not use ms-dos - netbsd-NetBSD. Do not use NetBSD or org.netbsd - netware-Novell NetWare. Do not use novell or NetWare - openbsd-OpenBSD. Do not use OpenBSD or org.openbsd - opensolaris-OpenSolaris. Do not use OpenSolaris or org.opensolaris - opensuse-openSUSE. Do not use suse, SuSE, or org.opensuse - rhel-Red Hat Enterprise Linux. Do not use redhat, RedHat, or com.redhat - sled-SUSE Linux Enterprise Desktop. Do not use com.suse - ubuntu-Ubuntu. Do not use Ubuntu, com.ubuntu, org.ubuntu, or canonical - windows-Microsoft Windows. Do not use com.microsoft.server
All	os_distro		

Specific to	Key	Description	Supported values
All	os_version	The operating system version as specified by the distributor.	Version number (for example, "11.10")
All	ramdisk_id	The ID of image stored in the Image Service that should be used as the ramdisk when booting an AMI-style image.	Valid image ID
All	vm_mode	The virtual machine mode. This represents the host/guest ABI (application binary interface) used for the virtual machine.	**hvm**-Fully virtualized. This is the mode used by QEMU and KVM.
libvirt API driver	hw_disk_bus	Specifies the type of disk controller to attach disk devices to.	**scsi**, **virtio**, **ide**, or **usb**. Note that if using **iscsi**, the **hw_scsi_model** needs to be set to **virtio-scsi**.
libvirt API driver	hw_numa_nodes	Number of NUMA nodes to expose to the instance (does not override flavor definition).	Integer. For a detailed example of NUMA-topology definition, see the hw:NUMA_def key in [Add Metadata]().
libvirt API driver	hw_numa_cpus.0	Mapping of vCPUs N-M to NUMA node 0 (does not override flavor definition).	Comma-separated list of integers.
libvirt API driver	hw_numa_cpus.1	Mapping of vCPUs N-M to NUMA node 1 (does not override flavor definition).	Comma-separated list of integers.

Specific to	Key	Description	Supported values
libvirt API driver	hw_numa_mem.0	Mapping N MB of RAM to NUMA node 0 (does not override flavor definition).	Integer
libvirt API driver	hw_numa_mem.1	Mapping N MB of RAM to NUMA node 1 (does not override flavor definition).	Integer
libvirt API driver	hw_qemu_guest_agent	Guest agent support. If set to **yes**, and if **qemu-ga** is also installed, file systems can be quiesced (frozen) and snapshots created automatically.	**yes** / **no**

Specific to	Key	Description	Supported values
libvirt API driver	hw_rng_model	Adds a random-number generator device to the image's instances. The cloud administrator can enable and control device behavior by configuring the instance's flavor. By default: • The generator device is disabled. • /dev/random is used as the default entropy source. To specify a physical HW RNG device, set **rng_dev_path** to "/dev/hwrng" in your Compute environment file.	**virtio**, or other supported device.

Specific to	Key	Description	Supported values
libvirt API driver	hw_scsi_model	Enables the use of VirtIO SCSI (virtio-scsi) to provide block device access for compute instances; by default, instances use VirtIO Block (virtio-blk). VirtIO SCSI is a para-virtualized SCSI controller device that provides improved scalability and performance, and supports advanced SCSI hardware.	**virtio-scsi**
libvirt API driver	hw_video_model	The video image driver used.	**vga**, **cirrus**, **vmvga**, **xen**, or **qxl**
libvirt API driver	hw_video_ram	Maximum RAM for the video image. Used only if a **hw_video:ram_max_mb** value has been set in the flavor's **extra_specs** and that value is higher than the value set in **hw_video_ram**.	Integer in MB (for example, *64*)

Specific to	Key	Description	Supported values
libvirt API driver	hw_watchdog_action	Enables a virtual hardware watchdog device that carries out the specified action if the server hangs. The watchdog uses the i6300esb device (emulating a PCI Intel 6300ESB). If **hw_watchdog_action** is not specified, the watchdog is disabled.	- disabled–The device is not attached. Allows the user to disable the watchdog for the image, even if it has been enabled using the image's flavor. The default value for this parameter is disabled. - reset–Forcefully reset the guest. - poweroff–Forcefully power off the guest. - pause–Pause the guest. - none–Only enable the watchdog; do nothing if the server hangs.
libvirt API driver	os_command_line	The kernel command line to be used by the libvirt driver, instead of the default. For Linux Containers (LXC), the value is used as arguments for initialization. This key is valid only for Amazon kernel, ramdisk, or machine images (aki, ari, or ami).	
libvirt API driver and VMware API driver	hw_vif_model	Specifies the model of virtual network interface device to use.	The valid options depend on the configured hypervisor. - KVM and QEMU: e1000, ne2k_pci, pcnet, rtl8139, and virtio. - VMware: e1000, e1000e, VirtualE1000, VirtualE1000e, VirtualPCNet32, VirtualSriovEthernetCard, and VirtualVmxnet. - Xen: e1000, netfront, ne2k_pci, pcnet, and rtl8139.

Specific to	Key	Description	Supported values
VMware API driver	vmware_adapter type	The virtual SCSI or IDE controller used by the hypervisor.	**lsiLogic**, **busLogic**, or **ide**
VMware API driver	vmware_ostype	A VMware GuestID which describes the operating system installed in the image. This value is passed to the hypervisor when creating a virtual machine. If not specified, the key defaults to **otherGuest**.	See thinkvirt.com.
VMware API driver	vmware_image_ version	Currently unused.	**1**
XenAPI driver	auto_disk_config	If true, the root partition on the disk is automatically resized before the instance boots. This value is only taken into account by the Compute service when using a Xen-based hypervisor with the XenAPI driver. The Compute service will only attempt to resize if there is a single partition on the image, and only if the partition is in **ext3** or **ext4** format.	**true** / **false**

Specific to	Key	Description	Supported values
libvirt API driver and XenAPI driver	os_type	The operating system installed on the image. The XenAPI driver contains logic that takes different actions depending on the value of the **os_type** parameter of the image. For example, for **os_type=windows** images, it creates a FAT32-based swap partition instead of a Linux swap partition, and it limits the injected host name to less than 16 characters.	**linux** or **windows**

Specific to	Key	Description	Supported values
libvirt API driver and XenAPI driver	os_type		

APPENDIX B. ENABLING THE LAUNCH INSTANCE WIZARD

There are two methods that you can use to launch instances from the dashboard:

- The Launch Instance form

- The Launch Instance wizard

The Launch Instance form is enabled by default, but you can enable the Launch Instance wizard at any time. You can also enable both the Launch Instance form and the Launch Instance wizard at the same time. The Launch Instance wizard simplifies the steps required to create instances.

1. Edit **/etc/openstack-dashboard/local_settings** file, and add the following values:

   ```
   LAUNCH_INSTANCE_LEGACY_ENABLED = False
   LAUNCH_INSTANCE_NG_ENABLED = True
   ```

2. Restart the *httpd* service:

   ```
   # systemctl restart httpd
   ```

The preferences for the Launch Instance form and Launch Instance wizard are updated.

If you enabled only one of these options, the **Launch Instance** button in the dashboard opens that option by default. If you enabled both options, two **Launch Instance** buttons are displayed in the dashboard, with the button on the left opening the Launch Instance wizard and the button on the right opening the Launch Instance form.