

OpenStack Installation Guide

for Red Hat Enterprise Linux, CentOS, and Fedora

icehouse (September 12, 2014)

Icehouse



OpenStack Installation Guide for Red Hat Enterprise Linux, CentOS, and Fedora

icehouse (2014-09-12)

Copyright © 2012, 2013 OpenStack Foundation All rights reserved.

The OpenStack® system consists of several key projects that you install separately but that work together depending on your cloud needs. These projects include Compute, Identity Service, Networking, Image Service, Block Storage, Object Storage, Telemetry, Orchestration, and Database. You can install any of these projects separately and configure them stand-alone or as connected entities. This guide shows you how to install OpenStack by using packages available through Fedora 20 as well as on Red Hat Enterprise Linux and its derivatives through the EPEL repository. Explanations of configuration options and sample configuration files are included.

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

Table of Contents

- Preface 7
 - Conventions 7
 - Document change history 7
- 1. Architecture 1
 - Overview 1
 - Conceptual architecture 2
 - Example architectures 3
- 2. Basic environment configuration 6
 - Before you begin 6
 - Networking 7
 - Network Time Protocol (NTP) 17
 - Security 17
 - Database 18
 - OpenStack packages 19
 - Messaging server 20
- 3. Configure the Identity Service 22
 - Identity Service concepts 22
 - Install the Identity Service 24
 - Define users, tenants, and roles 25
 - Define services and API endpoints 26
 - Verify the Identity Service installation 27
- 4. Install and configure the OpenStack clients 30
 - Overview 30
 - Install the OpenStack command-line clients 31
 - Set environment variables using the OpenStack RC file 33
 - Create openrc.sh files 34
- 5. Configure the Image Service 35
 - Image Service overview 35
 - Install the Image Service 36
 - Verify the Image Service installation 38
- 6. Configure Compute services 41
 - Compute service 41
 - Install Compute controller services 43
 - Configure a compute node 46
- 7. Add a networking service 48
 - OpenStack Networking (neutron) 48
 - Legacy networking (nova-network) 67
 - Next steps 69
- 8. Add the dashboard 70
 - System requirements 70
 - Install the dashboard 71
 - Next steps 72
- 9. Add the Block Storage service 73
 - Block Storage 73
 - Configure a Block Storage service controller 73
 - Configure a Block Storage service node 75
 - Verify the Block Storage installation 77
 - Next steps 78

10. Add Object Storage	79
Object Storage service	79
System requirements for Object Storage	80
Plan networking for Object Storage	80
Example of Object Storage installation architecture	82
Install Object Storage	83
Install and configure storage nodes	85
Install and configure the proxy node	86
Start services on the storage nodes	89
Verify the installation	89
Add another proxy server	90
Next steps	90
11. Add the Orchestration service	91
Orchestration service overview	91
Install the Orchestration service	91
Verify the Orchestration service installation	93
Next steps	95
12. Add the Telemetry module	96
Telemetry	96
Install the Telemetry module	97
Install the Compute agent for Telemetry	99
Configure the Image Service for Telemetry	100
Add the Block Storage service agent for Telemetry	101
Configure the Object Storage service for Telemetry	101
Verify the Telemetry installation	102
Next steps	103
13. Add the Database service	104
Database service overview	104
Install the Database service	105
Verify the Database service installation	108
14. Launch an instance	109
Launch an instance with OpenStack Networking (neutron)	109
Launch an instance with legacy networking (nova-network)	115
A. Reserved user IDs	121
B. Community support	122
Documentation	122
ask.openstack.org	123
OpenStack mailing lists	123
The OpenStack wiki	124
The Launchpad Bugs area	124
The OpenStack IRC channel	125
Documentation feedback	125
OpenStack distribution packages	125
Glossary	126

List of Tables

1.1. OpenStack services	1
2.1. Passwords	17
4.1. OpenStack services and clients	30
4.2. Prerequisite software	31
10.1. Hardware recommendations	80
A.1. Reserved user IDs	121

Preface

Conventions

The OpenStack documentation uses several typesetting conventions.

Notices

Notices take three forms:



Note

The information in a note is usually in the form of a handy tip or reminder.



Important

The information in an important notice is something you must be aware of before proceeding.



Warning

The information in warnings is critical. Warnings provide additional information about risk of data loss or security issues.

Command prompts

Commands prefixed with the # prompt are to be executed by the `root` user. These examples can also be executed by using the `sudo` command, if available.

Commands prefixed with the \$ prompt can be executed by any user, including `root`.

Document change history

This version of the guide replaces and obsoletes all previous versions. The following table describes the most recent changes:

Revision Date	Summary of Changes
April 16, 2014	<ul style="list-style-type: none">• Update for Icehouse, rework Networking setup to use ML2 as plugin, add new chapter for Database Service setup, improved basic configuration.
October 25, 2013	<ul style="list-style-type: none">• Added initial Debian support.
October 17, 2013	<ul style="list-style-type: none">• Havana release.
October 16, 2013	<ul style="list-style-type: none">• Add support for SUSE Linux Enterprise.
October 8, 2013	<ul style="list-style-type: none">• Complete reorganization for Havana.
September 9, 2013	<ul style="list-style-type: none">• Build also for openSUSE.
August 1, 2013	<ul style="list-style-type: none">• Fixes to Object Storage verification steps. Fix bug 1207347.
July 25, 2013	<ul style="list-style-type: none">• Adds creation of cinder user and addition to the service tenant. Fix bug 1205057.
May 8, 2013	<ul style="list-style-type: none">• Updated the book title for consistency.

1. Architecture

Table of Contents

Overview	1
Conceptual architecture	2
Example architectures	3

Overview

The *OpenStack* project is an open source cloud computing platform that supports all types of cloud environments. The project aims for simple implementation, massive scalability, and a rich set of features. Cloud computing experts from around the world contribute to the project.

OpenStack provides an Infrastructure-as-a-Service (*IaaS*) solution through a variety of complementary services. Each service offers an application programming interface (*API*) that facilitates this integration. The following table provides a list of OpenStack services:

Table 1.1. OpenStack services

Service	Project name	Description
<i>Dashboard</i>	<i>Horizon</i>	Provides a web-based self-service portal to interact with underlying OpenStack services, such as launching an instance, assigning IP addresses and configuring access controls.
<i>Compute</i>	<i>Nova</i>	Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning of virtual machines on demand.
<i>Networking</i>	<i>Neutron</i>	Enables network connectivity as a service for other OpenStack services, such as OpenStack Compute. Provides an API for users to define networks and the attachments into them. Has a pluggable architecture that supports many popular networking vendors and technologies.
Storage		
<i>Object Storage</i>	<i>Swift</i>	Stores and retrieves arbitrary unstructured data objects via a <i>RESTful</i> , HTTP based API. It is highly fault tolerant with its data replication and scale out architecture. Its implementation is not like a file server with mountable directories.
<i>Block Storage</i>	<i>Cinder</i>	Provides persistent block storage to running instances. Its pluggable driver architecture facilitates the creation and management of block storage devices.
Shared services		
<i>Identity service</i>	<i>Keystone</i>	Provides an authentication and authorization service for other OpenStack services. Provides a catalog of endpoints for all OpenStack services.
<i>Image Service</i>	<i>Glance</i>	Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.
<i>Telemetry</i>	<i>Ceilometer</i>	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability, and statistical purposes.

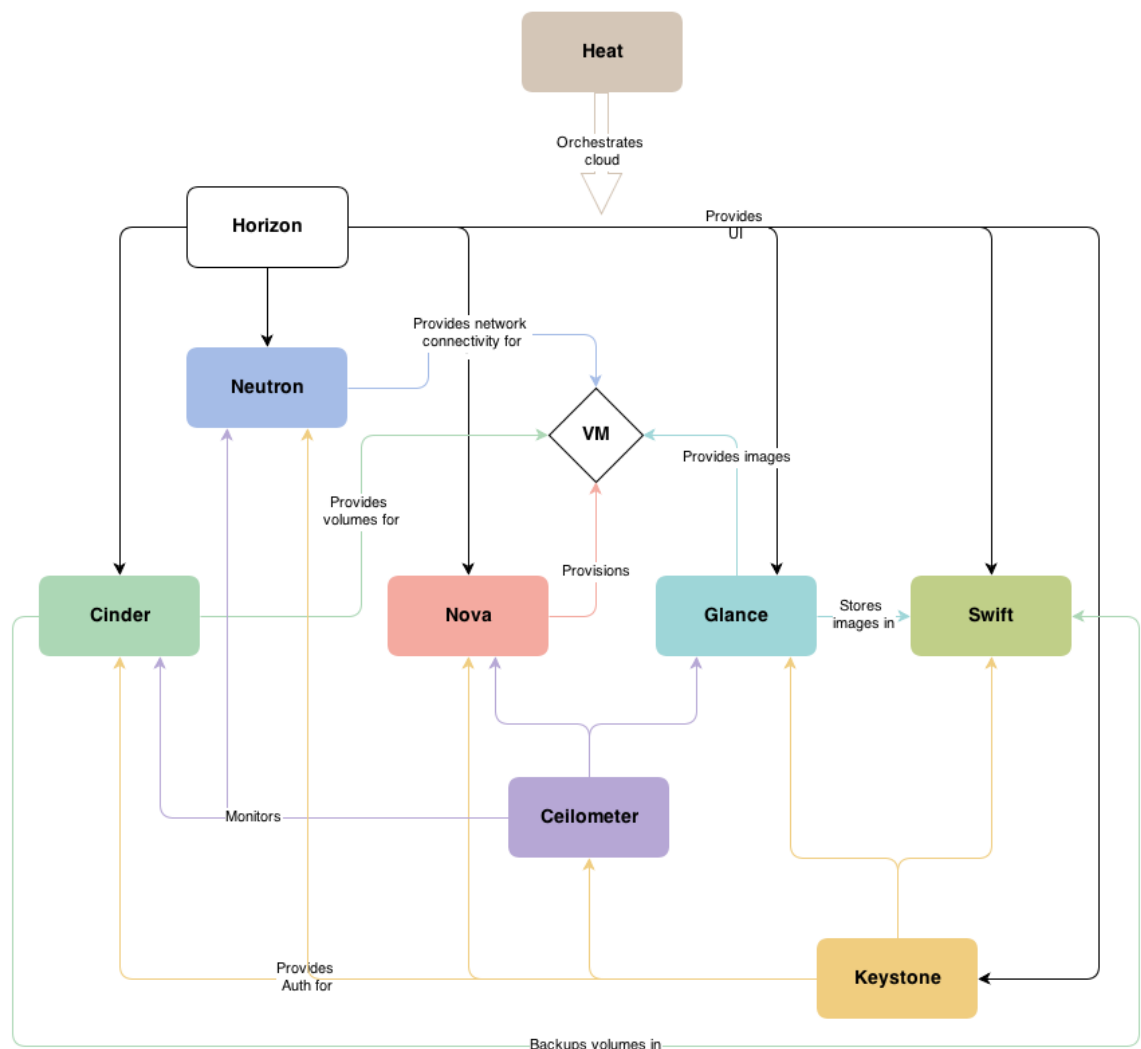
Service	Project name	Description
Higher-level services		
<i>Orchestration</i>	<i>Heat</i>	Orchestrates multiple composite cloud applications by using either the native <i>HOT</i> template format or the AWS CloudFormation template format, through both an OpenStack-native REST API and a CloudFormation-compatible Query API.
<i>Database Service</i>	<i>Trove</i>	Provides scalable and reliable Cloud Database-as-a-Service functionality for both relational and non-relational database engines.

This guide describes how to deploy these services in a functional test environment and, by example, teaches you how to build a production environment.

Conceptual architecture

Launching a virtual machine or instance involves many interactions among several services. The following diagram provides the conceptual architecture of a typical OpenStack environment.

Figure 1.1. Conceptual architecture



Example architectures

OpenStack is highly configurable to meet different needs with various compute, networking, and storage options. This guide enables you to choose your own OpenStack adventure using a combination of basic and optional services. This guide uses the following example architectures:

- Three-node architecture with OpenStack Networking (neutron). See [Figure 1.2, “Three-node architecture with OpenStack Networking \(neutron\)”](#) [4].
- The basic controller node runs the Identity service, Image Service, management portions of Compute and Networking, Networking plug-in, and the dashboard. It also includes supporting services such as a database, *message broker*, and *Network Time Protocol (NTP)*.

Optionally, the controller node also runs portions of Block Storage, Object Storage, Database Service, Orchestration, and Telemetry. These components provide additional features for your environment.

- The network node runs the Networking plug-in, layer 2 agent, and several layer 3 agents that provision and operate tenant networks. Layer 2 services include provisioning of virtual networks and tunnels. Layer 3 services include routing, *NAT*, and *DHCP*. This node also handles external (internet) connectivity for tenant virtual machines or instances.
- The compute node runs the hypervisor portion of Compute, which operates tenant virtual machines or instances. By default Compute uses KVM as the hypervisor. The compute node also runs the Networking plug-in and layer 2 agent which operate tenant networks and implement security groups. You can run more than one compute node.


Optionally, the compute node also runs the Telemetry agent. This component provides additional features for your environment.



Note

When you implement this architecture, skip [the section called “Legacy networking \(nova-network\)”](#) [67] in [Chapter 7, “Add a networking service”](#) [48]. To use optional services, you might need to install additional nodes, as described in subsequent chapters.

Figure 1.2. Three-node architecture with OpenStack Networking (neutron)



- Two-node architecture with legacy networking (nova-network). See [Figure 1.3, “Two-node architecture with legacy networking \(nova-network\)”](#) [5].
- The basic *controller node* runs the Identity service, Image Service, management portion of Compute, and the dashboard necessary to launch a simple instance. It also includes supporting services such as a database, message broker, and NTP.

Optionally, the controller node also runs portions of Block Storage, Object Storage, Database Service, Orchestration, and Telemetry. These components provide additional features for your environment.
- The basic *compute node* runs the *hypervisor* portion of Compute, which operates *tenant virtual machines* or instances. By default, Compute uses *KVM* as the *hypervisor*. Compute also provisions and operates tenant networks and implements *security groups*. You can run more than one compute node.

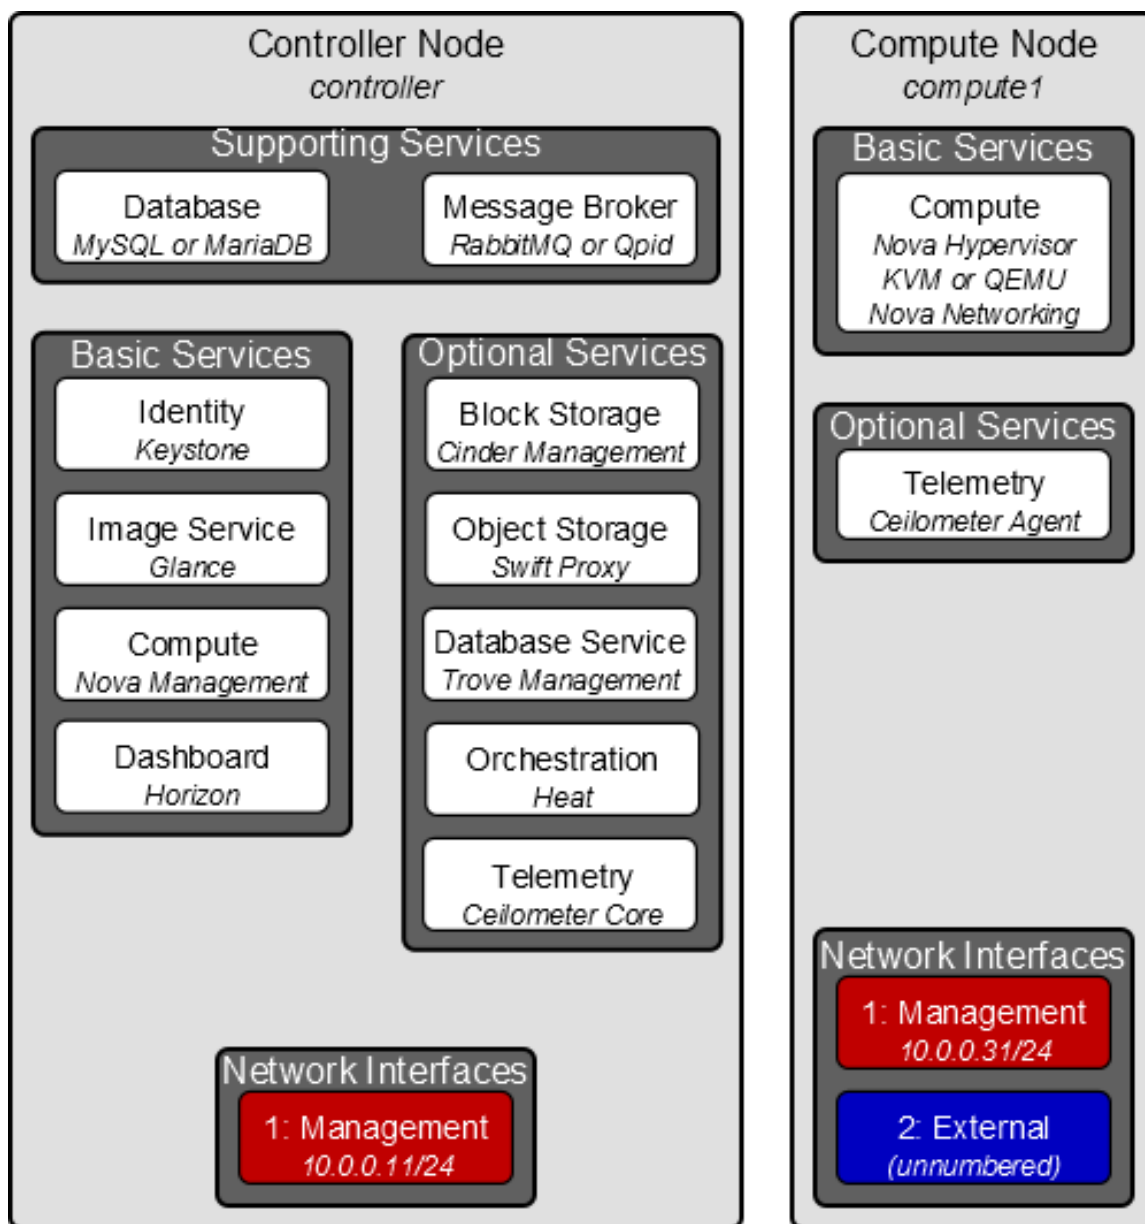
Optionally, the compute node also runs the Telemetry agent. This component provides additional features for your environment.



Note

When you implement this architecture, skip [the section called “OpenStack Networking \(neutron\)” \[48\]](#) in [Chapter 7, “Add a networking service” \[48\]](#). To use optional services, you might need to install additional nodes, as described in subsequent chapters.

Figure 1.3. Two-node architecture with legacy networking (nova-network)



2. Basic environment configuration

Table of Contents

Before you begin	6
Networking	7
Network Time Protocol (NTP)	17
Security	17
Database	18
OpenStack packages	19
Messaging server	20

This chapter explains how to configure each node in the [example architectures](#) including the [two-node architecture with legacy networking](#) and [three-node architecture with OpenStack Networking \(neutron\)](#).



Note

Although most environments include OpenStack Identity, Image Service, Compute, at least one networking service, and the dashboard, OpenStack Object Storage can operate independently of most other services. If your use case only involves Object Storage, you can skip to [the section called “System requirements for Object Storage” \[80\]](#). However, the dashboard will not work without at least OpenStack Image Service and Compute.



Note

You must use an account with administrative privileges to configure each node. Either run the commands as the `root` user or configure the `sudo` utility.

Before you begin

For a functional environment, OpenStack doesn't require a significant amount of resources. We recommend that your environment meets or exceeds the following minimum requirements which can support several minimal *CirrOS* instances:

- Controller Node: 1 processor, 2 GB memory, and 5 GB storage
- Network Node: 1 processor, 512 MB memory, and 5 GB storage
- Compute Node: 1 processor, 2 GB memory, and 10 GB storage

To minimize clutter and provide more resources for OpenStack, we recommend a minimal installation of your Linux distribution. Also, we strongly recommend that you install a 64-bit version of your distribution on at least the compute node. If you install a 32-bit version of your distribution on the compute node, attempting to start an instance using a 64-bit image will fail.



Note

A single disk partition on each node works for most basic installations. However, you should consider *Logical Volume Manager (LVM)* for installations with optional services such as Block Storage.

Many users build their test environments on *virtual machines (VMs)*. The primary benefits of VMs include the following:

- One physical server can support multiple nodes, each with almost any number of network interfaces.
- Ability to take periodic "snap shots" throughout the installation process and "roll back" to a working configuration in the event of a problem.

However, VMs will reduce performance of your instances, particularly if your hypervisor and/or processor lacks support for hardware acceleration of nested VMs.



Note

If you choose to install on VMs, make sure your hypervisor permits *promiscuous mode* on the *external network*.

For more information about system requirements, see the [OpenStack Operations Guide](#).

Networking

After installing the operating system on each node for the architecture that you choose to deploy, you must configure the network interfaces. We recommend that you disable any automated network management tools and manually edit the appropriate configuration files for your distribution. For more information on how to configure networking on your distribution, see the [documentation](#).

To disable NetworkManager and enable the network service:

- ```
service NetworkManager stop
service network start
chkconfig NetworkManager off
chkconfig network on
```

RHEL and derivatives including CentOS and Scientific Linux enable a restrictive *firewall* by default. During this installation, certain steps will fail unless you alter or disable the firewall. For further information about securing your installation, refer to the [OpenStack Security Guide](#).

On Fedora, `firewalld` replaces `iptables` as the default firewall system. While you can use `firewalld` successfully, this guide references `iptables` for compatibility with other distributions.

### To disable firewalld and enable iptables:

- ```
# service firewalld stop
# service iptables start
```

```
# chkconfig firewalld off
# chkconfig iptables on
```

Proceed to network configuration for the example [OpenStack Networking \(neutron\)](#) or [legacy networking \(nova-network\)](#) architecture.

OpenStack Networking (neutron)

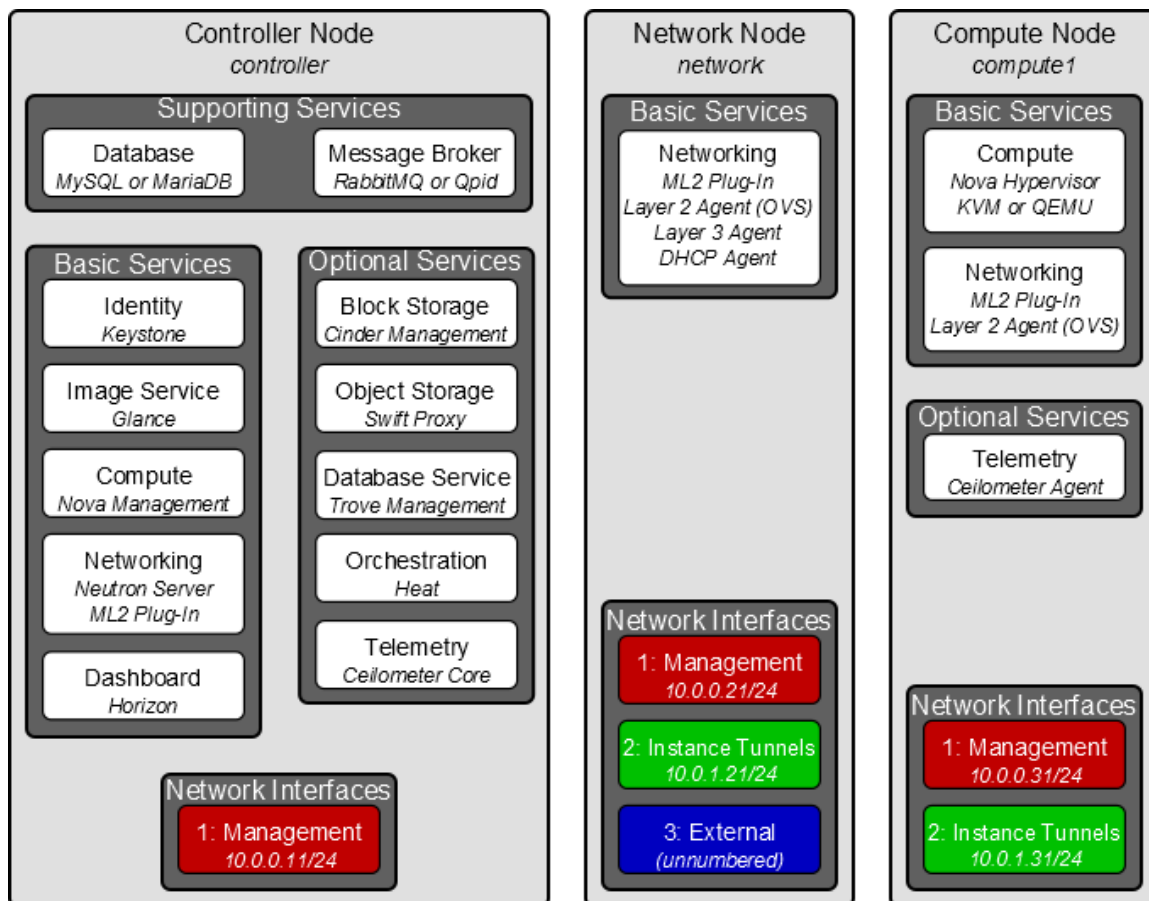
The example architecture with OpenStack Networking (neutron) requires one controller node, one network node, and at least one compute node. The controller node contains one network interface on the *management network*. The network node contains one network interface on the management network, one on the *instance tunnels network*, and one on the *external network*. The compute node contains one network interface on the management network and one on the instance tunnels network.



Note

Network interface names vary by distribution. Traditionally, interfaces use "eth" followed by a sequential number. To cover all variations, this guide simply refers to the first interface as the interface with the lowest number, the second interface as the interface with the middle number, and the third interface as the interface with the highest number.

Figure 2.1. Three-node architecture with OpenStack Networking (neutron)



Network mask: 255.255.255.0 (or /24)

3. The external interface uses a special configuration without an IP address assigned to it. Configure the third interface as the external interface:

Replace *INTERFACE_NAME* with the actual interface name. For example, *eth2* or *ens256*.

- Edit the `/etc/sysconfig/network-scripts/ifcfg-INTERFACE_NAME` file to contain the following:

Do not change the HWADDR and UUID keys.

```
DEVICE=INTERFACE_NAME
TYPE=Ethernet
ONBOOT="yes"
BOOTPROTO="none"
```

4. Restart networking:

```
# service network restart
```

To configure name resolution:

1. Set the hostname of the node to `network`.
2. Edit the `/etc/hosts` file to contain the following:

```
# network
10.0.0.21      network

# controller
10.0.0.11     controller

# compute1
10.0.0.31     compute1
```

Compute node

To configure networking:

1. Configure the first interface as the management interface:

IP address: 10.0.0.31

Network mask: 255.255.255.0 (or /24)

Default gateway: 10.0.0.1



Note

Additional compute nodes should use 10.0.0.32, 10.0.0.33, and so on.

2. Configure the second interface as the instance tunnels interface:

IP address: 10.0.1.31

Network mask: 255.255.255.0 (or /24)



Note

Additional compute nodes should use 10.0.1.32, 10.0.1.33, and so on.

To configure name resolution:

1. Set the hostname of the node to `computel`.
2. Edit the `/etc/hosts` file to contain the following:

```
# computel
10.0.0.31      computel

# controller
10.0.0.11     controller

# network
10.0.0.21     network
```

Verify connectivity

We recommend that you verify network connectivity to the internet and among the nodes before proceeding further.

1. From the *controller* node, **ping** a site on the internet:

```
# ping -c 4 openstack.org
PING openstack.org (174.143.194.225) 56(84) bytes of data.
64 bytes from 174.143.194.225: icmp_seq=1 ttl=54 time=18.3 ms
64 bytes from 174.143.194.225: icmp_seq=2 ttl=54 time=17.5 ms
64 bytes from 174.143.194.225: icmp_seq=3 ttl=54 time=17.5 ms
64 bytes from 174.143.194.225: icmp_seq=4 ttl=54 time=17.4 ms

--- openstack.org ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3022ms
rtt min/avg/max/mdev = 17.489/17.715/18.346/0.364 ms
```

2. From the *controller* node, **ping** the management interface on the *network* node:

```
# ping -c 4 network
PING network (10.0.0.21) 56(84) bytes of data.
64 bytes from network (10.0.0.21): icmp_seq=1 ttl=64 time=0.263 ms
64 bytes from network (10.0.0.21): icmp_seq=2 ttl=64 time=0.202 ms
64 bytes from network (10.0.0.21): icmp_seq=3 ttl=64 time=0.203 ms
64 bytes from network (10.0.0.21): icmp_seq=4 ttl=64 time=0.202 ms

--- network ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3000ms
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

3. From the *controller* node, **ping** the management interface on the *compute* node:

```
# ping -c 4 computel
```

```
PING compute1 (10.0.0.31) 56(84) bytes of data.  
64 bytes from compute1 (10.0.0.31): icmp_seq=1 ttl=64 time=0.263 ms  
64 bytes from compute1 (10.0.0.31): icmp_seq=2 ttl=64 time=0.202 ms  
64 bytes from compute1 (10.0.0.31): icmp_seq=3 ttl=64 time=0.203 ms  
64 bytes from compute1 (10.0.0.31): icmp_seq=4 ttl=64 time=0.202 ms  
  
--- network ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3000ms  
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

4. From the *network* node, ping a site on the internet:

```
# ping -c 4 openstack.org  
PING openstack.org (174.143.194.225) 56(84) bytes of data.  
64 bytes from 174.143.194.225: icmp_seq=1 ttl=54 time=18.3 ms  
64 bytes from 174.143.194.225: icmp_seq=2 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=3 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=4 ttl=54 time=17.4 ms  
  
--- openstack.org ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3022ms  
rtt min/avg/max/mdev = 17.489/17.715/18.346/0.364 ms
```

5. From the *network* node, ping the management interface on the *controller* node:

```
# ping -c 4 controller  
PING controller (10.0.0.11) 56(84) bytes of data.  
64 bytes from controller (10.0.0.11): icmp_seq=1 ttl=64 time=0.263 ms  
64 bytes from controller (10.0.0.11): icmp_seq=2 ttl=64 time=0.202 ms  
64 bytes from controller (10.0.0.11): icmp_seq=3 ttl=64 time=0.203 ms  
64 bytes from controller (10.0.0.11): icmp_seq=4 ttl=64 time=0.202 ms  
  
--- controller ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3000ms  
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

6. From the *network* node, ping the instance tunnels interface on the *compute* node:

```
# ping -c 4 10.0.1.31  
PING 10.0.1.31 (10.0.1.31) 56(84) bytes of data.  
64 bytes from 10.0.1.31 (10.0.1.31): icmp_seq=1 ttl=64 time=0.263 ms  
64 bytes from 10.0.1.31 (10.0.1.31): icmp_seq=2 ttl=64 time=0.202 ms  
64 bytes from 10.0.1.31 (10.0.1.31): icmp_seq=3 ttl=64 time=0.203 ms  
64 bytes from 10.0.1.31 (10.0.1.31): icmp_seq=4 ttl=64 time=0.202 ms  
  
--- 10.0.1.31 ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3000ms  
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

7. From the *compute* node, ping a site on the internet:

```
# ping -c 4 openstack.org  
PING openstack.org (174.143.194.225) 56(84) bytes of data.  
64 bytes from 174.143.194.225: icmp_seq=1 ttl=54 time=18.3 ms  
64 bytes from 174.143.194.225: icmp_seq=2 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=3 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=4 ttl=54 time=17.4 ms  
  
--- openstack.org ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3022ms
```

```
rtt min/avg/max/mdev = 17.489/17.715/18.346/0.364 ms
```

8. From the *compute* node, **ping** the management interface on the *controller* node:

```
# ping -c 4 controller
PING controller (10.0.0.11) 56(84) bytes of data.
64 bytes from controller (10.0.0.11): icmp_seq=1 ttl=64 time=0.263 ms
64 bytes from controller (10.0.0.11): icmp_seq=2 ttl=64 time=0.202 ms
64 bytes from controller (10.0.0.11): icmp_seq=3 ttl=64 time=0.203 ms
64 bytes from controller (10.0.0.11): icmp_seq=4 ttl=64 time=0.202 ms

--- controller ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3000ms
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

9. From the *compute* node, **ping** the instance tunnels interface on the *network* node:

```
# ping -c 4 10.0.1.21
PING 10.0.1.21 (10.0.1.21) 56(84) bytes of data.
64 bytes from 10.0.1.21 (10.0.1.21): icmp_seq=1 ttl=64 time=0.263 ms
64 bytes from 10.0.1.21 (10.0.1.21): icmp_seq=2 ttl=64 time=0.202 ms
64 bytes from 10.0.1.21 (10.0.1.21): icmp_seq=3 ttl=64 time=0.203 ms
64 bytes from 10.0.1.21 (10.0.1.21): icmp_seq=4 ttl=64 time=0.202 ms

--- 10.0.1.21 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3000ms
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

Legacy networking (nova-network)

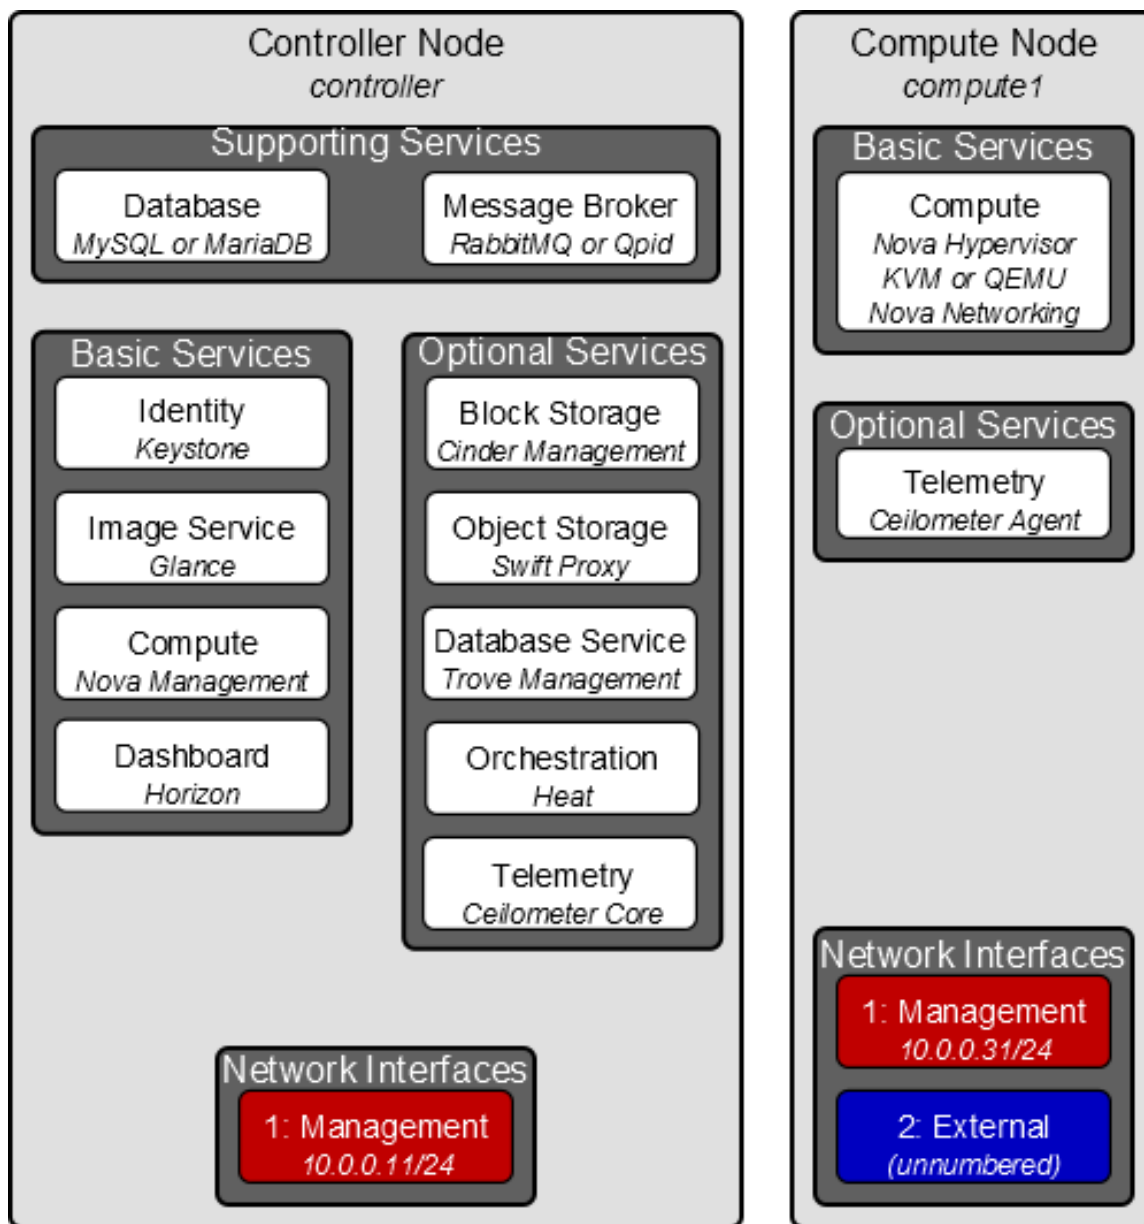
The example architecture with legacy networking (nova-network) requires a controller node and at least one compute node. The controller node contains one network interface on the *management network*. The compute node contains one network interface on the management network and one on the *external network*.



Note

Network interface names vary by distribution. Traditionally, interfaces use "eth" followed by a sequential number. To cover all variations, this guide simply refers to the first interface as the interface with the lowest number and the second interface as the interface with the highest number.

Figure 2.2. Two-node architecture with legacy networking (nova-network)



Unless you intend to use the exact configuration provided in this example architecture, you must modify the networks in this procedure to match your environment. Also, each node must resolve the other nodes by name in addition to IP address. For example, the *controller* name must resolve to 10.0.0.11, the IP address of the management interface on the controller node.



Warning

Reconfiguring network interfaces will interrupt network connectivity. We recommend using a local terminal session for these procedures.

Controller node

To configure networking:

- Configure the first interface as the management interface:

IP address: 10.0.0.11

Network mask: 255.255.255.0 (or /24)

Default gateway: 10.0.0.1

To configure name resolution:

1. Set the hostname of the node to *controller*.
2. Edit the `/etc/hosts` file to contain the following:

```
# controller
10.0.0.11      controller

# compute1
10.0.0.31     compute1
```

Compute node

To configure networking:

1. Configure the first interface as the management interface:

IP address: 10.0.0.31

Network mask: 255.255.255.0 (or /24)

Default gateway: 10.0.0.1



Note

Additional compute nodes should use 10.0.0.32, 10.0.0.33, and so on.

2. The external interface uses a special configuration without an IP address assigned to it. Configure the second interface as the external interface:

Replace *INTERFACE_NAME* with the actual interface name. For example, *eth1* or *ens224*.

- Edit the `/etc/sysconfig/network-scripts/ifcfg-INTERFACE_NAME` file to contain the following:

Do not change the HWADDR and UUID keys.

```
DEVICE=INTERFACE_NAME
TYPE=Ethernet
ONBOOT="yes"
BOOTPROTO="none"
```

-
- Restart networking:

```
# service network restart
```

To configure name resolution:

- Set the hostname of the node to `compute1`.
- Edit the `/etc/hosts` file to contain the following:

```
# compute1  
10.0.0.31      compute1  
  
# controller  
10.0.0.11     controller
```

Verify connectivity

We recommend that you verify network connectivity to the internet and among the nodes before proceeding further.

- From the *controller* node, ping a site on the internet:

```
# ping -c 4 openstack.org  
PING openstack.org (174.143.194.225) 56(84) bytes of data.  
64 bytes from 174.143.194.225: icmp_seq=1 ttl=54 time=18.3 ms  
64 bytes from 174.143.194.225: icmp_seq=2 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=3 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=4 ttl=54 time=17.4 ms  
  
--- openstack.org ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3022ms  
rtt min/avg/max/mdev = 17.489/17.715/18.346/0.364 ms
```

- From the *controller* node, ping the management interface on the *compute* node:

```
# ping -c 4 compute1  
PING compute1 (10.0.0.31) 56(84) bytes of data.  
64 bytes from compute1 (10.0.0.31): icmp_seq=1 ttl=64 time=0.263 ms  
64 bytes from compute1 (10.0.0.31): icmp_seq=2 ttl=64 time=0.202 ms  
64 bytes from compute1 (10.0.0.31): icmp_seq=3 ttl=64 time=0.203 ms  
64 bytes from compute1 (10.0.0.31): icmp_seq=4 ttl=64 time=0.202 ms  
  
--- compute1 ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3000ms  
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

- From the *compute* node, ping a site on the internet:

```
# ping -c 4 openstack.org  
PING openstack.org (174.143.194.225) 56(84) bytes of data.  
64 bytes from 174.143.194.225: icmp_seq=1 ttl=54 time=18.3 ms  
64 bytes from 174.143.194.225: icmp_seq=2 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=3 ttl=54 time=17.5 ms  
64 bytes from 174.143.194.225: icmp_seq=4 ttl=54 time=17.4 ms  
  
--- openstack.org ping statistics ---  
4 packets transmitted, 4 received, 0% packet loss, time 3022ms  
rtt min/avg/max/mdev = 17.489/17.715/18.346/0.364 ms
```


4. From the *compute* node, ping the management interface on the *controller* node:

```
# ping -c 4 controller
PING controller (10.0.0.11) 56(84) bytes of data.
64 bytes from controller (10.0.0.11): icmp_seq=1 ttl=64 time=0.263 ms
64 bytes from controller (10.0.0.11): icmp_seq=2 ttl=64 time=0.202 ms
64 bytes from controller (10.0.0.11): icmp_seq=3 ttl=64 time=0.203 ms
64 bytes from controller (10.0.0.11): icmp_seq=4 ttl=64 time=0.202 ms

--- controller ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3000ms
rtt min/avg/max/mdev = 0.202/0.217/0.263/0.030 ms
```

Network Time Protocol (NTP)

To synchronize services across multiple machines, you must install *NTP*. The examples in this guide configure the controller node as the reference server and any additional nodes to set their time from the controller node.

Install the `ntp` package on each system running OpenStack services:

```
# yum install ntp
```

Set up the NTP server on your controller node so that it receives data by modifying the `ntp.conf` file and restarting the service:

```
# service ntpd start
# chkconfig ntpd on
```

It is advised that you configure additional nodes to synchronize their time from the controller node rather than from outside of your LAN. To do so, install the `ntp` daemon as above, then edit `/etc/ntp.conf` and change the `server` directive to use the controller node as internet time source.

Security

OpenStack services support various security methods including password, policy, and encryption. Additionally, supporting services including the database server and message broker support at least password security.

To ease the installation process, this guide only covers password security where applicable. You can create secure passwords manually, generate them using a tool such as `pwgen`, or by running the following command:

```
$ openssl rand -hex 10
```

For OpenStack services, this guide uses `SERVICE_PASS` to reference service account passwords and `SERVICE_DBPASS` to reference database passwords.

The following table provides a list of services that require passwords and their associated references in the guide:

Table 2.1. Passwords

Password name	Description
Database password (no variable used)	Root password for the database

Password name	Description
<i>RABBIT_PASS</i>	Password of user guest of RabbitMQ
<i>KEYSTONE_DBPASS</i>	Database password of Identity service
<i>DEMO_PASS</i>	Password of user demo
<i>ADMIN_PASS</i>	Password of user admin
<i>GLANCE_DBPASS</i>	Database password for Image Service
<i>GLANCE_PASS</i>	Password of Image Service user glance
<i>NOVA_DBPASS</i>	Database password for Compute service
<i>NOVA_PASS</i>	Password of Compute service user nova
<i>DASH_DBPASS</i>	Database password for the dashboard
<i>CINDER_DBPASS</i>	Database password for the Block Storage service
<i>CINDER_PASS</i>	Password of Block Storage service user cinder
<i>NEUTRON_DBPASS</i>	Database password for the Networking service
<i>NEUTRON_PASS</i>	Password of Networking service user neutron
<i>HEAT_DBPASS</i>	Database password for the Orchestration service
<i>HEAT_PASS</i>	Password of Orchestration service user heat
<i>CEILOMETER_DBPASS</i>	Database password for the Telemetry service
<i>CEILOMETER_PASS</i>	Password of Telemetry service user ceilometer
<i>TROVE_DBPASS</i>	Database password of Database service
<i>TROVE_PASS</i>	Password of Database Service user trove

OpenStack and supporting services require administrative privileges during installation and operation. In some cases, services perform modifications to the host that can interfere with deployment automation tools such as Ansible, Chef, and Puppet. For example, some OpenStack services add a root wrapper to `sudo` that can interfere with security policies. See the [Cloud Administrator Guide](#) for more information. Also, the Networking service assumes default values for kernel network parameters and modifies firewall rules. To avoid most issues during your initial installation, we recommend using a stock deployment of a supported distribution on your hosts. However, if you choose to automate deployment of your hosts, review the configuration and policies applied to them before proceeding further.

Database

Most OpenStack services require a database to store information. These examples use a MySQL database that runs on the controller node. You must install the MySQL database on the controller node. You must install the MySQL Python library on any additional nodes that access MySQL.

Controller setup

On the controller node, install the MySQL client and server packages, and the Python library.

```
# yum install mysql mysql-server MySQL-python
```

The MySQL configuration requires some changes to work with OpenStack.

- Edit the `/etc/my.cnf` file:

- a. Under the `[mysqld]` section, set the `bind-address` key to the management IP address of the controller node to enable access by other nodes via the management network:

```
[mysqld]
...
bind-address = 10.0.0.11
```

- b. Under the `[mysqld]` section, set the following keys to enable InnoDB, UTF-8 character set, and UTF-8 collation by default:

```
[mysqld]
...
default-storage-engine = innodb
innodb_file_per_table
collation-server = utf8_general_ci
init-connect = 'SET NAMES utf8'
character-set-server = utf8
```

Start the MySQL database server and set it to start automatically when the system boots:

```
# service mysqld start
# chkconfig mysqld on
```

Finally, you should set a root password for your MySQL database. The OpenStack programs that set up databases and tables prompt you for this password if it is set.

You must delete the anonymous users that are created when the database is first started. Otherwise, database connection problems occur when you follow the instructions in this guide. To do this, use the `mysql_secure_installation` command. Note that if `mysql_secure_installation` fails you might need to use `mysql_install_db` first:

```
# mysql_install_db
# mysql_secure_installation
```

If you have not already set a root database password, press **ENTER** when you are prompted for the password. This command presents a number of options for you to secure your database installation. Respond **yes** to all prompts unless you have a good reason to do otherwise.

Node setup

On all nodes other than the controller node, install the MySQL Python library:

```
# yum install MySQL-python
```

OpenStack packages

Distributions might release OpenStack packages as part of their distribution or through other methods because the OpenStack and distribution release times are independent of each other.

This section describes the configuration you must complete after you configure machines to install the latest OpenStack packages.

The examples in this guide use the OpenStack packages from the RDO repository. These packages work on Red Hat Enterprise Linux 6, compatible versions of CentOS, and Fedora 20.

Install the yum-plugin-priorities plug-in. This package allows the assignment of relative priorities to the configured software repositories. This functionality is used by the RDO release packages:

```
# yum install yum-plugin-priorities
```

To enable the RDO repository, download and install the rdo-release-icehouse package:

```
# yum install http://repos.fedorapeople.org/repos/openstack/openstack-icehouse/rdo-release-icehouse-3.noarch.rpm
```

The EPEL package includes GPG keys for package signing and repository information. This should only be installed on Red Hat Enterprise Linux and CentOS, not Fedora. Install the latest epel-release package (see http://download.fedoraproject.org/pub/epel/6/x86_64/repoview/epel-release.html). For example:

```
# yum install http://dl.fedoraproject.org/pub/epel/6/x86_64/epel-release-6-8.noarch.rpm
```

The openstack-utils package contains utility programs that make installation and configuration easier. These programs are used throughout this guide. Install openstack-utils. This verifies that you can access the RDO repository:

```
# yum install openstack-utils
```



Warning

The openstack-config program in the openstack-utils package uses crudini to manipulate configuration files. However, crudini version 0.3 does not support multi valued options. See <https://bugs.launchpad.net/openstack-manuals/+bug/1269271>. As a work around, you must manually set any multi valued options or the new value overwrites the previous value instead of creating a new option.

The openstack-selinux package includes the policy files that are required to configure SELinux during OpenStack installation on RHEL and CentOS. This step is not required during OpenStack installation on Fedora. Install openstack-selinux:

```
# yum install openstack-selinux
```

Upgrade your system packages:

```
# yum upgrade
```

If the upgrade included a new kernel package, reboot the system to ensure the new kernel is running:

```
# reboot
```

Messaging server

OpenStack uses a *message broker* to coordinate operations and status information among services. The message broker service typically runs on the controller node. OpenStack

3. Configure the Identity Service

Table of Contents

Identity Service concepts	22
Install the Identity Service	24
Define users, tenants, and roles	25
Define services and API endpoints	26
Verify the Identity Service installation	27

Identity Service concepts

The *Identity Service* performs the following functions:

- **User management.** Tracks users and their permissions.
- **Service catalog.** Provides a catalog of available services with their API endpoints.

To understand the Identity Service, you must understand the following concepts:

User Digital representation of a person, system, or service who uses OpenStack cloud services. The Identity Service validates that incoming requests are made by the user who claims to be making the call. Users have a login and may be assigned tokens to access resources. Users can be directly assigned to a particular tenant and behave as if they are contained in that tenant.

Credentials Data that is known only by a user that proves who they are. In the Identity Service, examples are: User name and password, user name and API key, or an authentication token provided by the Identity Service.

Authentication The act of confirming the identity of a user. The Identity Service confirms an incoming request by validating a set of credentials supplied by the user.

These credentials are initially a user name and password or a user name and API key. In response to these credentials, the Identity Service issues an authentication token to the user, which the user provides in subsequent requests.

Token An arbitrary bit of text that is used to access resources. Each token has a scope which describes which resources are accessible with it. A token may be revoked at any time and is valid for a finite duration.

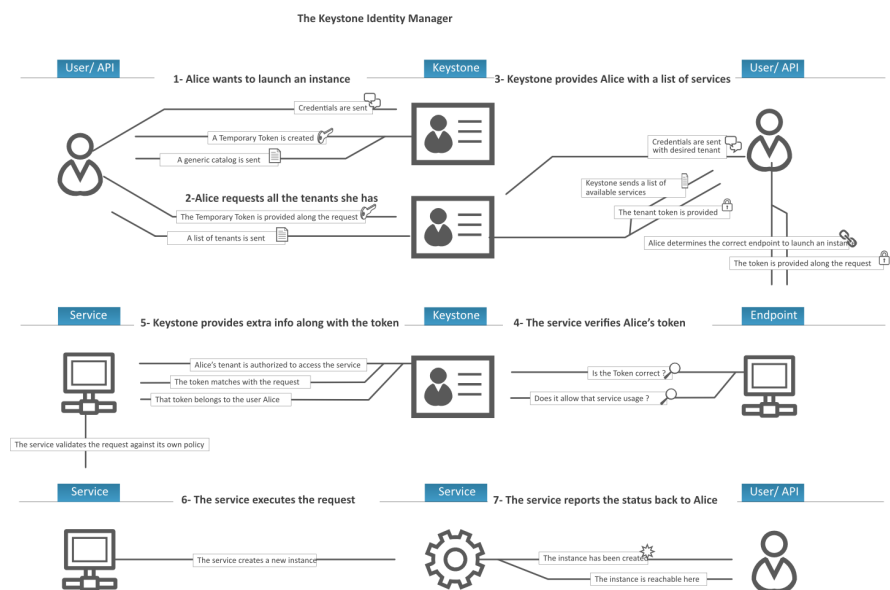
While the Identity Service supports token-based authentication in this release, the intention is for it to support additional protocols in the future. The

intent is for it to be an integration service foremost, and not aspire to be a full-fledged identity store and management solution.

<i>Tenant</i>	A container used to group or isolate resources and/or identity objects. Depending on the service operator, a tenant may map to a customer, account, organization, or project.
<i>Service</i>	An OpenStack service, such as Compute (Nova), Object Storage (Swift), or Image Service (Glance). Provides one or more endpoints through which users can access resources and perform operations.
<i>Endpoint</i>	A network-accessible address, usually described by a URL, from where you access a service. If using an extension for templates, you can create an endpoint template, which represents the templates of all the consumable services that are available across the regions.
<i>Role</i>	A personality that a user assumes that enables them to perform a specific set of operations. A role includes a set of rights and privileges. A user assuming that role inherits those rights and privileges.

In the Identity Service, a token that is issued to a user includes the list of roles that user has. Services that are being called by that user determine how they interpret the set of roles a user has and to which operations or resources each role grants access.

The following diagram shows the Identity Service process flow:



Install the Identity Service

1. Install the OpenStack Identity Service on the controller node, together with `python-keystoneclient` (which is a dependency):

```
# yum install openstack-keystone python-keystoneclient
```

2. The Identity Service uses a database to store information. Specify the location of the database in the configuration file. In this guide, we use a MySQL database on the controller node with the username `keystone`. Replace `KEYSTONE_DBPASS` with a suitable password for the database user.

```
# openstack-config --set /etc/keystone/keystone.conf \
    database connection mysql://keystone:KEYSTONE_DBPASS@controller/
keystone
```

3. Use the password that you set previously to log in as root. Create a `keystone` database user:

```
$ mysql -u root -p
mysql> CREATE DATABASE keystone;
mysql> GRANT ALL PRIVILEGES ON keystone.* TO 'keystone'@'localhost' \
    IDENTIFIED BY 'KEYSTONE_DBPASS';
mysql> GRANT ALL PRIVILEGES ON keystone.* TO 'keystone'@'% ' \
    IDENTIFIED BY 'KEYSTONE_DBPASS';
mysql> exit
```

4. Create the database tables for the Identity Service:

```
# su -s /bin/sh -c "keystone-manage db_sync" keystone
```

5. Define an authorization token to use as a shared secret between the Identity Service and other OpenStack services. Use `openssl` to generate a random token and store it in the configuration file:

```
# ADMIN_TOKEN=$(openssl rand -hex 10)
# echo $ADMIN_TOKEN
# openstack-config --set /etc/keystone/keystone.conf DEFAULT \
    admin_token $ADMIN_TOKEN
```

6. By default, Keystone uses PKI tokens. Create the signing keys and certificates and restrict access to the generated data:

```
# keystone-manage pki_setup --keystone-user keystone --keystone-group
keystone
# chown -R keystone:keystone /etc/keystone/ssl
# chmod -R o-rwx /etc/keystone/ssl
```

7. Start the Identity Service and enable it to start when the system boots:

```
# service openstack-keystone start
# chkconfig openstack-keystone on
```

8. By default, the Identity Service stores expired tokens in the database indefinitely. While potentially useful for auditing in production environments, the accumulation of expired tokens will considerably increase database size and may decrease service performance, particularly in test environments with limited resources. We recommend configuring a periodic task using `cron` to purge expired tokens hourly.

- Run the following command to purge expired tokens every hour and log the output to `/var/log/keystone/keystone-tokenflush.log`:

```
# (crontab -l -u keystone 2>&1 | grep -q token_flush) || \
echo '@hourly /usr/bin/keystone-manage token_flush >/var/log/keystone/
keystone-tokenflush.log 2>&1' >> /var/spool/cron/keystone
```

Define users, tenants, and roles

After you install the Identity Service, set up *users*, *tenants*, and *roles* to authenticate against. These are used to allow access to services and *endpoints*, described in the next section.

Typically, you would indicate a user and password to authenticate with the Identity Service. At this point, however, you have not created any users, so you have to use the authorization token created in an earlier step, see [the section called “Install the Identity Service” \[24\]](#) for further details. You can pass this with the `--os-token` option to the `keystone` command or set the `OS_SERVICE_TOKEN` environment variable. Set `OS_SERVICE_TOKEN`, as well as `OS_SERVICE_ENDPOINT` to specify where the Identity Service is running. Replace `ADMIN_TOKEN` with your authorization token.

```
$ export OS_SERVICE_TOKEN=ADMIN_TOKEN
$ export OS_SERVICE_ENDPOINT=http://controller:35357/v2.0
```

Create an administrative user

Follow these steps to create an administrative user, role, and tenant. You will use this account for administrative interaction with the OpenStack cloud.

By default, the Identity Service creates a special `_member_` role. The OpenStack dashboard automatically grants access to users with this role. You will give the `admin` user access to this role in addition to the `admin` role.



Note

Any role that you create must map to roles specified in the `policy.json` file included with each OpenStack service. The default policy file for most services grants administrative access to the `admin` role.

1. Create the `admin` user:

```
$ keystone user-create --name=admin --pass=ADMIN_PASS --email=ADMIN_EMAIL
```

Replace `ADMIN_PASS` with a secure password and replace `ADMIN_EMAIL` with an email address to associate with the account.

2. Create the `admin` role:

```
$ keystone role-create --name=admin
```

3. Create the `admin` tenant:

```
$ keystone tenant-create --name=admin --description="Admin Tenant"
```

4. You must now link the admin user, admin role, and admin tenant together using the `user-role-add` option:

```
$ keystone user-role-add --user=admin --tenant=admin --role=admin
```

5. Link the admin user, `_member_` role, and admin tenant:

```
$ keystone user-role-add --user=admin --role=_member_ --tenant=admin
```

Create a normal user

Follow these steps to create a normal user and tenant, and link them to the special `_member_` role. You will use this account for daily non-administrative interaction with the OpenStack cloud. You can also repeat this procedure to create additional cloud users with different usernames and passwords. Skip the tenant creation step when creating these users.

1. Create the demo user:

```
$ keystone user-create --name=demo --pass=DEMO_PASS --email=DEMO_EMAIL
```

Replace `DEMO_PASS` with a secure password and replace `DEMO_EMAIL` with an email address to associate with the account.

2. Create the demo tenant:

```
$ keystone tenant-create --name=demo --description="Demo Tenant"
```



Note

Do not repeat this step when adding additional users.

3. Link the demo user, `_member_` role, and demo tenant:

```
$ keystone user-role-add --user=demo --role=_member_ --tenant=demo
```

Create a service tenant

OpenStack services also require a username, tenant, and role to access other OpenStack services. In a basic installation, OpenStack services typically share a single tenant named `service`.

You will create additional usernames and roles under this tenant as you install and configure each service.

- Create the `service` tenant:

```
$ keystone tenant-create --name=service --description="Service Tenant"
```

Define services and API endpoints

So that the Identity Service can track which OpenStack services are installed and where they are located on the network, you must register each service in your OpenStack installation. To register a service, run these commands:

- **keystone service-create**. Describes the service.
- **keystone endpoint-create**. Associates *API endpoints* with the service.

You must also register the Identity Service itself. Use the `OS_SERVICE_TOKEN` environment variable, as set previously, for authentication.

1. Create a service entry for the Identity Service:

```
$ keystone service-create --name=keystone --type=identity \  
--description="OpenStack Identity"
+-----+-----+
| Property | Value |
+-----+-----+
| description | OpenStack Identity |
| id | 15c11a23667e427e91bc31335b45f4bd |
| name | keystone |
| type | identity |
+-----+-----+
```

The service ID is randomly generated and is different from the one shown here.

2. Specify an API endpoint for the Identity Service by using the returned service ID. When you specify an endpoint, you provide URLs for the public API, internal API, and admin API. In this guide, the `controller` host name is used. Note that the Identity Service uses a different port for the admin API.

```
$ keystone endpoint-create \  
--service-id=$(keystone service-list | awk '/ identity / {print $2}') \  
--publicurl=http://controller:5000/v2.0 \  
--internalurl=http://controller:5000/v2.0 \  
--adminurl=http://controller:35357/v2.0
+-----+-----+
| Property | Value |
+-----+-----+
| adminurl | http://controller:35357/v2.0 |
| id | 11f9c625a3b94a3f8e66bf4e5de2679f |
| internalurl | http://controller:5000/v2.0 |
| publicurl | http://controller:5000/v2.0 |
| region | regionOne |
| service_id | 15c11a23667e427e91bc31335b45f4bd |
+-----+-----+
```



Note

You will need to create an additional endpoint for each service added to your OpenStack environment. The sections of this guide associated with the installation of each service include the endpoint creation step specific to the service.

Verify the Identity Service installation

1. To verify that the Identity Service is installed and configured correctly, clear the values in the `OS_SERVICE_TOKEN` and `OS_SERVICE_ENDPOINT` environment variables:

```
$ unset OS_SERVICE_TOKEN OS_SERVICE_ENDPOINT
```

These variables, which were used to bootstrap the administrative user and register the Identity Service, are no longer needed.

2. You can now use regular user name-based authentication.

Request a authentication token by using the `admin` user and the password you chose for that user:

```
$ keystone --os-username=admin --os-password=ADMIN_PASS \
  --os-auth-url=http://controller:35357/v2.0 token-get
```

In response, you receive a token paired with your user ID. This verifies that the Identity Service is running on the expected endpoint and that your user account is established with the expected credentials.

3. Verify that authorization behaves as expected. To do so, request authorization on a tenant:

```
$ keystone --os-username=admin --os-password=ADMIN_PASS \
  --os-tenant-name=admin --os-auth-url=http://controller:35357/v2.0 \
  token-get
```

In response, you receive a token that includes the ID of the tenant that you specified. This verifies that your user account has an explicitly defined role on the specified tenant and the tenant exists as expected.

4. You can also set your `--os-*` variables in your environment to simplify command-line usage. Set up a `admin-openrc.sh` file with the admin credentials and admin endpoint:

```
export OS_USERNAME=admin
export OS_PASSWORD=ADMIN_PASS
export OS_TENANT_NAME=admin
export OS_AUTH_URL=http://controller:35357/v2.0
```

5. Source this file to read in the environment variables:

```
$ source admin-openrc.sh
```

6. Verify that your `admin-openrc.sh` file is configured correctly. Run the same command without the `--os-*` arguments:

```
$ keystone token-get
```

The command returns a token and the ID of the specified tenant. This verifies that you have configured your environment variables correctly.

7. Verify that your admin account has authorization to perform administrative commands:

```
$ keystone user-list
```

id	name	enabled	email
afea5bde3be9413dbd60e479fddf9228	admin	True	admin@example.com
32aca1f9a47540c29d6988091f76c934	demo	True	demo@example.com

```
$ keystone user-role-list --user admin --tenant admin
+-----+
+-----+-----+-----+
|          id          | name |          user_id          |
|          | tenant_id |          |
+-----+-----+-----+
| 9fe2ff9ee4384b1894a90878d3e92bab | _member_ | e519b772cb43474582fa303da62559e5 |
| afea5bde3be9413dbd60e479fddf9228 | admin | e519b772cb43474582fa303da62559e5 |
| 5d3b60b66f1f438b80eaae41a77b5951 |          | e519b772cb43474582fa303da62559e5 |
| afea5bde3be9413dbd60e479fddf9228 |          | e519b772cb43474582fa303da62559e5 |
+-----+-----+-----+
```

Seeing that the `id` in the output from the `keystone user-list` command matches the `user_id` in the `keystone user-role-list` command, and that the `admin` role is listed for that user, for the related tenant, this verifies that your user account has the `admin` role, which matches the role used in the Identity Service `policy.json` file.



Note

As long as you define your credentials and the Identity Service endpoint through the command line or environment variables, you can run all OpenStack client commands from any machine. For details, see [Chapter 4, “Install and configure the OpenStack clients”](#) [30].

Prerequisite	Description
	<pre># apt-get install python-pip</pre>
	<p>Red Hat Enterprise Linux, CentOS, or Fedora. A packaged version available in RDO enables you to use yum to install the clients, or you can install pip and use it to manage client installation:</p>
	<pre># yum install python-pip</pre>
	<p>openSUSE 12.2 and earlier. A packaged version available in the Open Build Service enables you to use rpm or zypper to install the clients, or you can install pip and use it to manage client installation:</p>
	<pre># zypper install python-pip</pre>
	<p>openSUSE 12.3 and later. A packaged version enables you to use rpm or zypper to install the clients. See the section called "Install the clients" [32]</p>

Install the clients

When following the instructions in this section, replace *PROJECT* with the lowercase name of the client to install, such as **nova**. Repeat for each client. The following values are valid:

- `ceilometer` - Telemetry API
- `cinder` - Block Storage API and extensions
- `glance` - Image Service API
- `heat` - Orchestration API
- `keystone` - Identity service API and extensions
- `neutron` - Networking API
- `nova` - Compute API and extensions
- `swift` - Object Storage API
- `trove` - Database Service API

The following example shows the command for installing the nova client with *pip*.

```
# pip install python-novaclient
```

Installing with pip

Use pip to install the OpenStack clients on a Linux, Mac OS X, or Microsoft Windows system. It is easy to use and ensures that you get the latest version of the client from the [Python Package Index](#). Also, pip enables you to update or remove a package.

Install each client separately by using the following command:

- For Mac OS X or Linux:

```
# pip install python-PROJECTclient
```

- For Microsoft Windows:


```
C:\>pip install python-PROJECTclient
```

Installing from packages

RDO and openSUSE have client packages that can be installed without `pip`.

On Red Hat Enterprise Linux, CentOS, or Fedora, use `yum` to install the clients from the packaged versions available in [RDO](#):

```
# yum install python-PROJECTclient
```

For openSUSE, use `rpm` or `zypper` to install the clients from the packaged versions available in [the Open Build Service](#):

```
# zypper install python-PROJECT
```

Upgrade or remove clients

To upgrade a client, add the `--upgrade` option to the `pip install` command:

```
# pip install --upgrade python-PROJECTclient
```

To remove the a client, run the `pip uninstall` command:

```
# pip uninstall python-PROJECTclient
```

Set environment variables using the OpenStack RC file

To set the required environment variables for the OpenStack command-line clients, you must create an environment file called an OpenStack rc file, or `openrc.sh` file. This project-specific environment file contains the credentials that all OpenStack services use.

When you source the file, environment variables are set for your current shell. The variables enable the OpenStack client commands to communicate with the OpenStack services that run in the cloud.



Note

Defining environment variables using an environment file is not a common practice on Microsoft Windows. Environment variables are usually defined in the **Advanced** tab of the System Properties dialog box.

Create and source the OpenStack RC file

1. In a text editor, create a file named `PROJECT-openrc.sh` file and add the following authentication information:

The following example shows the information for a project called `admin`, where the OS username is also `admin`, and the identity host is located at `controller`.

```
export OS_USERNAME=admin
export OS_PASSWORD=ADMIN_PASS
export OS_TENANT_NAME=admin
export OS_AUTH_URL=http://controller:35357/v2.0
```

2. On any shell from which you want to run OpenStack commands, source the `PROJECT-openrc.sh` file for the respective project. In this example, you source the `admin-openrc.sh` file for the `admin` project:

```
$ source admin-openrc.sh
```

Override environment variable values

When you run OpenStack client commands, you can override some environment variable settings by using the options that are listed at the end of the **help** output of the various client commands. For example, you can override the `OS_PASSWORD` setting in the `PROJECT-openrc.sh` file by specifying a password on a **keystone** command, as follows:

```
$ keystone --os-password PASSWORD service-list
```

Where `PASSWORD` is your password.

Create `openrc.sh` files

As explained in [the section called “Create and source the OpenStack RC file” \[33\]](#), use the credentials from [the section called “Define users, tenants, and roles” \[25\]](#) and create the following `PROJECT-openrc.sh` files:

- `admin-openrc.sh` for the administrative user
- `demo-openrc.sh` for the normal user:

```
export OS_USERNAME=demo
export OS_PASSWORD=DEMO_PASS
export OS_TENANT_NAME=demo
export OS_AUTH_URL=http://controller:35357/v2.0
```

5. Configure the Image Service

Table of Contents

Image Service overview	35
Install the Image Service	36
Verify the Image Service installation	38

The OpenStack Image Service enables users to discover, register, and retrieve virtual machine images. Also known as the glance project, the Image Service offers a *REST* API that enables you to query virtual machine image metadata and retrieve an actual image. You can store virtual machine images made available through the Image Service in a variety of locations from simple file systems to object-storage systems like OpenStack Object Storage.



Important

For simplicity, this guide configures the Image Service to use the `file` back end. This means that images uploaded to the Image Service are stored in a directory on the same system that hosts the service. By default, this directory is `/var/lib/glance/images/`.

Before you proceed, ensure that the system has sufficient space available in this directory to store virtual machine images and snapshots. At an absolute minimum, several gigabytes of space should be available for use by the Image Service in a proof of concept deployment. To see requirements for other back ends, see [Configuration Reference](#).

Image Service overview

The Image Service includes the following components:

- `glance-api`. Accepts Image API calls for image discovery, retrieval, and storage.
- `glance-registry`. Stores, processes, and retrieves metadata about images. Metadata includes items such as size and type.



Security note

The registry is a private internal service meant only for use by the Image Service itself. Do not expose it to users.

- Database. Stores image metadata. You can choose your database depending on your preference. Most deployments use MySQL or SQLite.
- Storage repository for image files. The Image Service supports a variety of repositories including normal file systems, Object Storage, RADOS block devices, HTTP, and Amazon S3. Some types of repositories support only read-only usage.

A number of periodic processes run on the Image Service to support caching. Replication services ensures consistency and availability through the cluster. Other periodic processes include auditors, updaters, and reapers.

As shown in [Figure 1.1, “Conceptual architecture” \[2\]](#), the Image Service is central to the overall IaaS picture. It accepts API requests for images or image metadata from end users or Compute components and can store its disk files in the Object Storage Service.

Install the Image Service

The OpenStack Image Service acts as a registry for virtual disk images. Users can add new images or take a snapshot of an image from an existing server for immediate storage. Use snapshots for back up and as templates to launch new servers. You can store registered images in Object Storage or in other locations. For example, you can store images in simple file systems or external web servers.



Note

This procedure assumes you set the appropriate environment variables to your credentials as described in [the section called “Verify the Identity Service installation” \[27\]](#).

1. Install the Image Service on the controller node:

```
# yum install openstack-glance python-glanceclient
```

2. The Image Service stores information about images in a database. The examples in this guide use the MySQL database that is used by other OpenStack services.

Configure the location of the database. The Image Service provides the `glance-api` and `glance-registry` services, each with its own configuration file. You must update both configuration files throughout this section. Replace `GLANCE_DBPASS` with your Image Service database password.

```
# openstack-config --set /etc/glance/glance-api.conf database \
connection mysql://glance:GLANCE_DBPASS@controller/glance
# openstack-config --set /etc/glance/glance-registry.conf database \
connection mysql://glance:GLANCE_DBPASS@controller/glance
```

3. Use the password you created to log in as root and create a `glance` database user:

```
$ mysql -u root -p
mysql> CREATE DATABASE glance;
mysql> GRANT ALL PRIVILEGES ON glance.* TO 'glance'@'localhost' \
IDENTIFIED BY 'GLANCE_DBPASS';
mysql> GRANT ALL PRIVILEGES ON glance.* TO 'glance'@'%' \
IDENTIFIED BY 'GLANCE_DBPASS';
```

4. Create the database tables for the Image Service:

```
# su -s /bin/sh -c "glance-manage db_sync" glance
```

5. Create a `glance` user that the Image Service can use to authenticate with the Identity service. Choose a password and specify an email address for the `glance` user. Use the service tenant and give the user the `admin` role:

```
$ keystone user-create --name=glance --pass=GLANCE_PASS \  
  --email=glance@example.com  
$ keystone user-role-add --user=glance --tenant=service --role=admin
```

6. Configure the Image Service to use the Identity Service for authentication.

Run the following commands and replace `GLANCE_PASS` with the password you chose for the `glance` user in the Identity Service:

```
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
auth_uri http://controller:5000  
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
auth_host controller  
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
auth_port 35357  
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
auth_protocol http  
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
admin_tenant_name service  
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
admin_user glance  
# openstack-config --set /etc/glance/glance-api.conf keystone_authtoken \  
admin_password GLANCE_PASS  
# openstack-config --set /etc/glance/glance-api.conf paste_deploy \  
flavor keystone  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
auth_uri http://controller:5000  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
auth_host controller  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
auth_port 35357  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
auth_protocol http  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
admin_tenant_name service  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
admin_user glance  
# openstack-config --set /etc/glance/glance-registry.conf  
keystone_authtoken \  
admin_password GLANCE_PASS  
# openstack-config --set /etc/glance/glance-registry.conf paste_deploy \  
flavor keystone
```

7. Register the Image Service with the Identity service so that other OpenStack services can locate it. Register the service and create the endpoint:

```
$ keystone service-create --name=glance --type=image \  
  --description="OpenStack Image Service"  
$ keystone endpoint-create \  
  --service-id=$(keystone service-list | awk '/ image / {print $2}') \  
  --publicurl=http://controller:9292 \  
  --internalurl=http://controller:9292 \  
  --adminurl=http://controller:9292
```

8. Start the `glance-api` and `glance-registry` services and configure them to start when the system boots:

```
# service openstack-glance-api start
# service openstack-glance-registry start
# chkconfig openstack-glance-api on
# chkconfig openstack-glance-registry on
```

Verify the Image Service installation

To test the Image Service installation, download at least one virtual machine image that is known to work with OpenStack. For example, CirrOS is a small test image that is often used for testing OpenStack deployments ([CirrOS downloads](#)). This walk through uses the 64-bit CirrOS QCOW2 image.

For more information about how to download and build images, see [OpenStack Virtual Machine Image Guide](#). For information about how to manage images, see the [OpenStack User Guide](#).

1. Download the image into a dedicated directory using `wget` or `curl`:

```
$ mkdir /tmp/images
$ cd /tmp/images/
$ wget http://cdn.download.cirros-cloud.net/0.3.2/cirros-0.3.2-x86_64-disk.img
```

2. Upload the image to the Image Service:

```
$ glance image-create --name=IMAGELABEL --disk-format=FILEFORMAT \
--container-format=CONTAINERFORMAT --is-public=ACCESSVALUE < IMAGEFILE
```

Where:

IMAGELABEL Arbitrary label. The name by which users refer to the image.

FILEFORMAT Specifies the format of the image file. Valid formats include `qcow2`, `raw`, `vhd`, `vmdk`, `vdi`, `iso`, `aki`, `ari`, and `ami`.

You can verify the format using the `file` command:

```
$ file cirros-0.3.2-x86_64-disk.img
cirros-0.3.2-x86_64-disk.img: QEMU QCOW Image (v2),
41126400 bytes
```

CONTAINERFORMAT Specifies the container format. Valid formats include: `bare`, `ovf`, `aki`, `ari` and `ami`.

Specify `bare` to indicate that the image file is not in a file format that contains metadata about the virtual machine. Although this field is currently required, it is not actually used by any of the OpenStack services and has no effect on system behavior. Because the value is not used anywhere, it is safe to always specify `bare` as the container format.

ACCESSVALUE Specifies image access:

- true - All users can view and use the image.
- false - Only administrators can view and use the image.

IMAGEFILE Specifies the name of your downloaded image file.

For example:

```
$ source admin-openrc.sh
$ glance image-create --name "cirros-0.3.2-x86_64" --disk-format qcow2 \
  --container-format bare --is-public True --progress < cirros-0.3.2-
  x86_64-disk.img
```

Property	Value
checksum	64d7c1cd2b6f60c92c14662941cb7913
container_format	bare
created_at	2014-04-08T18:59:18
deleted	False
deleted_at	None
disk_format	qcow2
id	acafc7c0-40aa-4026-9673-b879898e1fc2
is_public	True
min_disk	0
min_ram	0
name	cirros-0.3.2-x86_64
owner	efa984b0a914450e9a47788ad330699d
protected	False
size	13167616
status	active
updated_at	2014-01-08T18:59:18



Note

Because the returned image ID is generated dynamically, your deployment generates a different ID than the one shown in this example.

- 3. Confirm that the image was uploaded and display its attributes:

```
$ glance image-list
```

ID	Name	Disk Format
Container Format	Size	Status
acafc7c0-40aa-4026-9673-b879898e1fc2	cirros-0.3.2-x86_64	qcow2
bare	13167616	active

- 4. You can now remove the locally downloaded image, since it is stored and available through the Image Service.

```
$ rm -r /tmp/images
```


6. Configure Compute services

Table of Contents

Compute service	41
Install Compute controller services	43
Configure a compute node	46

Compute service

The Compute service is a cloud computing fabric controller, which is the main part of an IaaS system. Use it to host and manage cloud computing systems. The main modules are implemented in Python.

Compute interacts with the Identity Service for authentication, Image Service for images, and the Dashboard for the user and administrative interface. Access to images is limited by project and by user; quotas are limited per project (for example, the number of instances). The Compute service scales horizontally on standard hardware, and downloads images to launch instances as required.

The Compute service is made up of the following functional areas and their underlying components:

API

- `nova-api` service. Accepts and responds to end user compute API calls. Supports the OpenStack Compute API, the Amazon EC2 API, and a special Admin API for privileged users to perform administrative actions. Also, initiates most orchestration activities, such as running an instance, and enforces some policies.
- `nova-api-metadata` service. Accepts metadata requests from instances. The `nova-api-metadata` service is generally only used when you run in multi-host mode with `nova-network` installations. For details, see [Metadata service](#) in the *Cloud Administrator Guide*.

On Debian systems, it is included in the `nova-api` package, and can be selected through `debconf`.

Compute core

- `nova-compute` process. A worker daemon that creates and terminates virtual machine instances through hypervisor APIs. For example, XenAPI for XenServer/XCP, libvirt for KVM or QEMU, VMwareAPI for VMware, and so on. The process by which it does so is fairly complex but the basics are simple: Accept actions from the queue and perform a series of system commands, like launching a KVM instance, to carry them out while updating state in the database.
- `nova-scheduler` process. Conceptually the simplest piece of code in Compute. Takes a virtual machine instance request from the queue and determines on which compute server host it should run.

Console interface

- `nova-consoleauth` daemon. Authorizes tokens for users that console proxies provide. See `nova-novncproxy` and `nova-xvncproxy`. This service must be running for console proxies to work. Many proxies of either type can be run against a single `nova-consoleauth` service in a cluster configuration. For information, see [About nova-consoleauth](#).
- `nova-novncproxy` daemon. Provides a proxy for accessing running instances through a VNC connection. Supports browser-based novnc clients.
- `nova-xvncproxy` daemon. A proxy for accessing running instances through a VNC connection. Supports a Java client specifically designed for OpenStack.
- `nova-cert` daemon. Manages x509 certificates.

Image management (EC2 scenario)

- `nova-objectstore` daemon. Provides an S3 interface for registering images with the Image Service. Mainly used for installations that must support euca2ools. The euca2ools tools talk to `nova-objectstore` in *S3 language*, and `nova-objectstore` translates S3 requests into Image Service requests.
- euca2ools client. A set of command-line interpreter commands for managing cloud resources. Though not an OpenStack module, you can configure `nova-api` to support this EC2 interface. For more information, see the [Eucalyptus 3.4 Documentation](#).

Command-line clients and other interfaces

- nova client. Enables users to submit commands as a tenant administrator or end user.
- nova-manage client. Enables cloud administrators to submit commands.

Other components

- The queue. A central hub for passing messages between daemons. Usually implemented with [RabbitMQ](#), but could be any AMQP message queue, such as [Apache Qpid](#) or [Zero MQ](#).
- SQL database. Stores most build-time and runtime states for a cloud infrastructure. Includes instance types that are available for use, instances in use, available networks, and projects. Theoretically, OpenStack Compute can support any database that SQLAlchemy supports, but the only databases widely used are SQLite3 databases (only appropriate for test and development work), MySQL, and PostgreSQL.

The Compute service interacts with other OpenStack services: Identity Service for authentication, Image Service for images, and the OpenStack dashboard for a web interface.

Install Compute controller services

Compute is a collection of services that enable you to launch virtual machine instances. You can configure these services to run on separate nodes or the same node. In this guide, most

services run on the controller node and the service that launches virtual machines runs on a dedicated compute node. This section shows you how to install and configure these services on the controller node.

1. Install the Compute packages necessary for the controller node.

```
# yum install openstack-nova-api openstack-nova-cert openstack-nova-conductor \
    openstack-nova-console openstack-nova-novncproxy openstack-nova-scheduler \
    python-novaclient
```

2. Compute stores information in a database. In this guide, we use a MySQL database on the controller node. Configure Compute with the database location and credentials. Replace `NOVA_DBPASS` with the password for the database that you will create in a later step.

```
# openstack-config --set /etc/nova/nova.conf \
    database_connection mysql://nova:NOVA_DBPASS@controller/nova
```

3. Set these configuration keys to configure Compute to use the Qpid message broker:

```
# openstack-config --set /etc/nova/nova.conf \
    DEFAULT rpc_backend qpid
# openstack-config --set /etc/nova/nova.conf DEFAULT
    qpid_hostname controller
```

4. Set the `my_ip`, `vncserver_listen`, and `vncserver_proxyclient_address` configuration options to the management interface IP address of the controller node:

```
# openstack-config --set /etc/nova/nova.conf DEFAULT my_ip 10.0.0.11
# openstack-config --set /etc/nova/nova.conf DEFAULT vncserver_listen 10.0.0.11
# openstack-config --set /etc/nova/nova.conf DEFAULT
    vncserver_proxyclient_address 10.0.0.11
```

5. Use the password you created previously to log in as root. Create a nova database user:

```
$ mysql -u root -p
mysql> CREATE DATABASE nova;
mysql> GRANT ALL PRIVILEGES ON nova.* TO 'nova'@'localhost' \
    IDENTIFIED BY 'NOVA_DBPASS';
mysql> GRANT ALL PRIVILEGES ON nova.* TO 'nova'@ '%' \
    IDENTIFIED BY 'NOVA_DBPASS';
```

6. Create the Compute service tables:

```
# su -s /bin/sh -c "nova-manage db sync" nova
```

7. Create a nova user that Compute uses to authenticate with the Identity Service. Use the service tenant and give the user the admin role:

```
$ keystone user-create --name=nova --pass=NOVA_PASS --email=nova@example.com
$ keystone user-role-add --user=nova --tenant=service --role=admin
```

8. Configure Compute to use these credentials with the Identity Service running on the controller. Replace `NOVA_PASS` with your Compute password.

```
# openstack-config --set /etc/nova/nova.conf DEFAULT auth_strategy
keystone
# openstack-config --set /etc/nova/nova.conf keystone_auth_token auth_uri
http://controller:5000
# openstack-config --set /etc/nova/nova.conf keystone_auth_token
auth_host controller
# openstack-config --set /etc/nova/nova.conf keystone_auth_token
auth_protocol http
# openstack-config --set /etc/nova/nova.conf keystone_auth_token auth_port
35357
# openstack-config --set /etc/nova/nova.conf keystone_auth_token admin_user
nova
# openstack-config --set /etc/nova/nova.conf keystone_auth_token
admin_tenant_name service
# openstack-config --set /etc/nova/nova.conf keystone_auth_token
admin_password NOVA_PASS
```

9. You must register Compute with the Identity Service so that other OpenStack services can locate it. Register the service and specify the endpoint:

```
$ keystone service-create --name=nova --type=compute \
--description="OpenStack Compute"
$ keystone endpoint-create \
--service-id=$(keystone service-list | awk '/ compute / {print $2}') \
--publicurl=http://controller:8774/v2/%(tenant_id)s \
--internalurl=http://controller:8774/v2/%(tenant_id)s \
--adminurl=http://controller:8774/v2/%(tenant_id)s
```

10. Start Compute services and configure them to start when the system boots:

```
# service openstack-nova-api start
# service openstack-nova-cert start
# service openstack-nova-consoleauth start
# service openstack-nova-scheduler start
# service openstack-nova-conductor start
# service openstack-nova-novncproxy start
# chkconfig openstack-nova-api on
# chkconfig openstack-nova-cert on
# chkconfig openstack-nova-consoleauth on
# chkconfig openstack-nova-scheduler on
# chkconfig openstack-nova-conductor on
# chkconfig openstack-nova-novncproxy on
```

11. To verify your configuration, list available images:

```
$ nova image-list
+-----+-----+-----+
+-----+
| ID                               | Name                               | Status |
| Server |
+-----+-----+-----+
| acafc7c0-40aa-4026-9673-b879898e1fc2 | cirros-0.3.2-x86_64 | ACTIVE |
| |
+-----+-----+-----+
+-----+
```

Configure a compute node

After you configure the Compute service on the controller node, you must configure another system as a compute node. The compute node receives requests from the controller node and hosts virtual machine instances. You can run all services on a single node, but the examples in this guide use separate systems. This makes it easy to scale horizontally by adding additional Compute nodes following the instructions in this section.

The Compute service relies on a hypervisor to run virtual machine instances. OpenStack can use various hypervisors, but this guide uses KVM.

1. Install the Compute packages:

```
# yum install openstack-nova-compute
```

2. Edit the `/etc/nova/nova.conf` configuration file:

```
# openstack-config --set /etc/nova/nova.conf database connection mysql://  
nova:NOVA_DBPASS@controller/nova  
# openstack-config --set /etc/nova/nova.conf DEFAULT auth_strategy  
keystone  
# openstack-config --set /etc/nova/nova.conf keystone_auth token auth_uri  
http://controller:5000  
# openstack-config --set /etc/nova/nova.conf keystone_auth token  
auth_host controller  
# openstack-config --set /etc/nova/nova.conf keystone_auth token  
auth_protocol http  
# openstack-config --set /etc/nova/nova.conf keystone_auth token auth_port  
35357  
# openstack-config --set /etc/nova/nova.conf keystone_auth token admin_user  
nova  
# openstack-config --set /etc/nova/nova.conf keystone_auth token  
admin_tenant_name service  
# openstack-config --set /etc/nova/nova.conf keystone_auth token  
admin_password NOVA_PASS
```

3. Configure the Compute service to use the Qpid message broker by setting these configuration keys:

```
# openstack-config --set /etc/nova/nova.conf \  
DEFAULT rpc_backend qpid  
# openstack-config --set /etc/nova/nova.conf DEFAULT  
qpid_hostname controller
```

4. Configure Compute to provide remote console access to instances.

```
# openstack-config --set /etc/nova/nova.conf DEFAULT my_ip 10.0.0.31  
# openstack-config --set /etc/nova/nova.conf DEFAULT vnc_enabled True  
# openstack-config --set /etc/nova/nova.conf DEFAULT vncserver_listen 0.0.  
0.0  
# openstack-config --set /etc/nova/nova.conf DEFAULT  
vncserver_proxyclient_address 10.0.0.31  
# openstack-config --set /etc/nova/nova.conf \  
DEFAULT novncproxy_base_url http://controller:6080/vnc_auto.html
```

5. Specify the host that runs the Image Service.

```
# openstack-config --set /etc/nova/nova.conf DEFAULT  
glance_host controller
```

-
6. You must determine whether your system's processor and/or hypervisor support hardware acceleration for virtual machines.

Run the following command:

```
$ egrep -c '(vmx|svm)' /proc/cpuinfo
```

If this command returns a value of *one or greater*, your system supports hardware acceleration which typically requires no additional configuration.

If this command returns a value of *zero*, your system does not support hardware acceleration and you must configure `libvirt` to use QEMU instead of KVM.

- Run the following command:

```
# openstack-config --set /etc/nova/nova.conf libvirt virt_type qemu
```

7. Start the Compute service and its dependencies. Configure them to start automatically when the system boots.

- For RHEL or CentOS:

```
# service libvirtd start
# service messagebus start
# service openstack-nova-compute start
# chkconfig libvirtd on
# chkconfig messagebus on
# chkconfig openstack-nova-compute on
```

- For Fedora:

```
# service libvirtd start
# service dbus start
# service openstack-nova-compute start
# chkconfig libvirtd on
# chkconfig dbus on
# chkconfig openstack-nova-compute on
```



Note

On RHEL 7.x and derivatives, the `dbus/messagebus` commands listed above may cause an error. These errors can be ignored because these services are now enabled by default.

You can associate external network IP addresses with ports to VMs. This way, entities on the outside network can access VMs.

Networking also supports *security groups*. Security groups enable administrators to define firewall rules in groups. A VM can belong to one or more security groups, and Networking applies the rules in those security groups to block or unblock ports, port ranges, or traffic types for that VM.

Each plug-in that Networking uses has its own concepts. While not vital to operating Networking, understanding these concepts can help you set up Networking. All Networking installations use a core plug-in and a security group plug-in (or just the No-Op security group plug-in). Additionally, Firewall-as-a-service (FWaaS) and Load-balancing-as-a-service (LBaaS) plug-ins are available.

Modular Layer 2 (ML2) plug-in

Configure controller node

Prerequisites

Before you configure OpenStack Networking (neutron), you must create a database and Identity service credentials including a user and service.

1. Connect to the database as the root user, create the `neutron` database, and grant the proper access to it:

Replace `NEUTRON_DBPASS` with a suitable password.

```
$ mysql -u root -p
mysql> CREATE DATABASE neutron;
mysql> GRANT ALL PRIVILEGES ON neutron.* TO 'neutron'@'localhost' \
IDENTIFIED BY 'NEUTRON_DBPASS';
mysql> GRANT ALL PRIVILEGES ON neutron.* TO 'neutron'@'%' \
IDENTIFIED BY 'NEUTRON_DBPASS';
```

2. Create Identity service credentials for Networking:

- a. Create the `neutron` user:

Replace `NEUTRON_PASS` with a suitable password and `neutron@example.com` with a suitable e-mail address.

```
$ keystone user-create --name neutron --pass NEUTRON_PASS --
email neutron@example.com
```

- b. Link the `neutron` user to the `service` tenant and `admin` role:

```
$ keystone user-role-add --user neutron --tenant service --role admin
```

- c. Create the `neutron` service:

```
$ keystone service-create --name neutron --type network --description
"OpenStack Networking"
```

- d. Create the service endpoint:

```
$ keystone endpoint-create \
```

```
--service-id $(keystone service-list | awk '/ network / {print $2}') \
--publicurl http://controller:9696 \
--adminurl http://controller:9696 \
--internalurl http://controller:9696
```

To install the Networking components

- ```
yum install openstack-neutron openstack-neutron-ml2 python-neutronclient
```

## To configure the Networking server component

The Networking server component configuration includes the database, authentication mechanism, message broker, topology change notifier, and plug-in.

1. Configure Networking to use the database:

Replace `NEUTRON_DBPASS` with a suitable password.

```
openstack-config --set /etc/neutron/neutron.conf database connection \
mysql://neutron:NEUTRON_DBPASS@controller/neutron
```

2. Configure Networking to use the Identity service for authentication:

Replace `NEUTRON_PASS` with the password you chose for the `neutron` user in the Identity service.

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
auth_strategy keystone
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
auth_uri http://controller:5000
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
auth_host controller
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
auth_protocol http
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
auth_port 35357
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
admin_tenant_name service
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
admin_user neutron
openstack-config --set /etc/neutron/neutron.conf keystone_auth token \
admin_password NEUTRON_PASS
```

3. Configure Networking to use the message broker:

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
rpc_backend neutron.openstack.common.rpc.impl_qpid
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
qpid_hostname controller
```

4. Configure Networking to notify Compute about network topology changes:

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
notify_nova_on_port_status_changes True
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
notify_nova_on_port_data_changes True
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
nova_url http://controller:8774/v2
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
```



```
network_api_class nova.network.neutronv2.api.API
openstack-config --set /etc/nova/nova.conf DEFAULT \
neutron_url http://controller:9696
openstack-config --set /etc/nova/nova.conf DEFAULT \
neutron_auth_strategy keystone
openstack-config --set /etc/nova/nova.conf DEFAULT \
neutron_admin_tenant_name service
openstack-config --set /etc/nova/nova.conf DEFAULT \
neutron_admin_username neutron
openstack-config --set /etc/nova/nova.conf DEFAULT \
neutron_admin_password NEUTRON_PASS
openstack-config --set /etc/nova/nova.conf DEFAULT \
neutron_admin_auth_url http://controller:35357/v2.0
openstack-config --set /etc/nova/nova.conf DEFAULT \
linuxnet_interface_driver nova.network.linux_net.LinuxOVSIInterfaceDriver
openstack-config --set /etc/nova/nova.conf DEFAULT \
firewall_driver nova.virt.firewall.NoopFirewallDriver
openstack-config --set /etc/nova/nova.conf DEFAULT \
security_group_api neutron
```



### Note

By default, Compute uses an internal firewall service. Since Networking includes a firewall service, you must disable the Compute firewall service by using the `nova.virt.firewall.NoopFirewallDriver` firewall driver.

## To finalize installation

1. The Networking service initialization scripts expect a symbolic link `/etc/neutron/plugin.ini` pointing to the configuration file associated with your chosen plug-in. Using ML2, for example, the symbolic link must point to `/etc/neutron/plugins/ml2/ml2_conf.ini`. If this symbolic link does not exist, create it using the following commands:

```
ln -s plugins/ml2/ml2_conf.ini /etc/neutron/plugin.ini
```

2. Restart the Compute services:

```
service openstack-nova-api restart
service openstack-nova-scheduler restart
service openstack-nova-conductor restart
```

3. Start the Networking service and configure it to start when the system boots:

```
service neutron-server start
chkconfig neutron-server on
```



### Note

Unlike other services, Networking typically does not require a separate step to populate the database because the `neutron-server` service populates it automatically. However, the packages for these distributions sometimes require running the `neutron-db-manage` command prior to starting the `neutron-server` service. We recommend attempting to start the service before manually populating the database. If the service returns database errors, perform the following operations:

1. Configure Networking to use long plug-in names:

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
 core_plugin neutron.plugins.ml2.plugin.Ml2Plugin
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
 service_plugins neutron.services.l3_router.
 l3_router_plugin.L3RouterPlugin
```

2. Populate the database:

```
su -s /bin/sh -c "neutron-db-manage --config-file /etc/
neutron/neutron.conf \
 --config-file /etc/neutron/plugin.ini upgrade head"
neutron
```

3. Attempt to start the `neutron-server` service again. You can return the `core_plugin` and `service_plugins` configuration keys to short plug-in names.

## Configure network node

### Prerequisites

Before you configure OpenStack Networking, you must enable certain kernel networking functions.

1. Edit `/etc/sysctl.conf` to contain the following:

```
net.ipv4.ip_forward=1
net.ipv4.conf.all.rp_filter=0
net.ipv4.conf.default.rp_filter=0
```

2. Implement the changes:

```
sysctl -p
```

### To install the Networking components

- ```
# yum install openstack-neutron openstack-neutron-ml2 \
  openstack-neutron-openvswitch
```

To configure the Networking common components

The Networking common component configuration includes the authentication mechanism, message broker, and plug-in.

1. Configure Networking to use the Identity service for authentication:

Replace `NEUTRON_PASS` with the password you chose for the `neutron` user in the Identity service.

```
# openstack-config --set /etc/neutron/neutron.conf DEFAULT \
  auth_strategy keystone
# openstack-config --set /etc/neutron/neutron.conf keystone_authtoken \
  auth_uri http://controller:5000
# openstack-config --set /etc/neutron/neutron.conf keystone_authtoken \
  auth_host controller
# openstack-config --set /etc/neutron/neutron.conf keystone_authtoken \
```

```
auth_protocol http
# openstack-config --set /etc/neutron/neutron.conf keystone_auth_token \
auth_port 35357
# openstack-config --set /etc/neutron/neutron.conf keystone_auth_token \
admin_tenant_name service
# openstack-config --set /etc/neutron/neutron.conf keystone_auth_token \
admin_user neutron
# openstack-config --set /etc/neutron/neutron.conf keystone_auth_token \
admin_password NEUTRON_PASS
```

2. Configure Networking to use the message broker:

```
# openstack-config --set /etc/neutron/neutron.conf DEFAULT \
rpc_backend neutron.openstack.common.rpc.impl_qpid
# openstack-config --set /etc/neutron/neutron.conf DEFAULT \
qpid_hostname controller
```

3. Configure Networking to use the Modular Layer 2 (ML2) plug-in and associated services:

```
# openstack-config --set /etc/neutron/neutron.conf DEFAULT \
core_plugin ml2
# openstack-config --set /etc/neutron/neutron.conf DEFAULT \
service_plugins router
```



Note

We recommend adding `verbose = True` to the `[DEFAULT]` section in `/etc/neutron/neutron.conf` to assist with troubleshooting.

To configure the Layer-3 (L3) agent

The *Layer-3 (L3) agent* provides routing services for instance virtual networks.

- Run the following commands:

```
# openstack-config --set /etc/neutron/l3_agent.ini DEFAULT \
interface_driver neutron.agent.linux.interface.OVSInterfaceDriver
# openstack-config --set /etc/neutron/l3_agent.ini DEFAULT \
use_namespaces True
```



Note

We recommend adding `verbose = True` to the `[DEFAULT]` section in `/etc/neutron/l3_agent.ini` to assist with troubleshooting.

To configure the DHCP agent

The *DHCP agent* provides *DHCP* services for instance virtual networks.

1. Run the following commands:

```
# openstack-config --set /etc/neutron/dhcp_agent.ini DEFAULT \
interface_driver neutron.agent.linux.interface.OVSInterfaceDriver
# openstack-config --set /etc/neutron/dhcp_agent.ini DEFAULT \
dhcp_driver neutron.agent.linux.dhcp.Dnsmasq
# openstack-config --set /etc/neutron/dhcp_agent.ini DEFAULT \
use_namespaces True
```

**Note**

We recommend adding `verbose = True` to the `[DEFAULT]` section in `/etc/neutron/dhcp_agent.ini` to assist with troubleshooting.

2. Tunneling protocols such as *generic routing encapsulation (GRE)* include additional packet headers that increase overhead and decrease space available for the payload or user data. Without knowledge of the virtual network infrastructure, instances attempt to send packets using the default Ethernet *maximum transmission unit (MTU)* of 1500 bytes. *Internet protocol (IP)* networks contain the *path MTU discovery (PMTUD)* mechanism to detect end-to-end MTU and adjust packet size accordingly. However, some operating systems and networks block or otherwise lack support for PMTUD causing performance degradation or connectivity failure.

Ideally, you can prevent these problems by enabling *jumbo frames* on the physical network that contains your tenant virtual networks. Jumbo frames support MTUs up to approximately 9000 bytes which negates the impact of GRE overhead on virtual networks. However, many network devices lack support for jumbo frames and OpenStack administrators often lack control of network infrastructure. Given the latter complications, you can also prevent MTU problems by reducing the instance MTU to account for GRE overhead. Determining the proper MTU value often takes experimentation, but 1454 bytes works in most environments. You can configure the *DHCP* server that assigns IP addresses to your instances to also adjust the MTU.

**Note**

Some cloud images such as CirrOS ignore the DHCP MTU option.

- a. Run the following command:

```
# openstack-config --set /etc/neutron/dhcp_agent.ini DEFAULT \
  dnsmasq_config_file /etc/neutron/dnsmasq-neutron.conf
```

- b. Create and edit the `/etc/neutron/dnsmasq-neutron.conf` file and add the following keys:

```
dhcp-option-force=26,1454
```

- c. Kill any existing `dnsmasq` processes:

```
# killall dnsmasq
```

To configure the metadata agent

The *metadata agent* provides configuration information such as credentials for remote access to instances.

1. Run the following commands:

Replace `NEUTRON_PASS` with the password you chose for the `neutron` user in the Identity service. Replace `METADATA_SECRET` with a suitable secret for the metadata proxy.

```
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
```

```

auth_url http://controller:5000/v2.0
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
  auth_region regionOne
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
  admin_tenant_name service
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
  admin_user neutron
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
  admin_password NEUTRON_PASS
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
  nova_metadata_ip controller
# openstack-config --set /etc/neutron/metadata_agent.ini DEFAULT \
  metadata_proxy_shared_secret METADATA_SECRET

```



Note

We recommend adding `verbose = True` to the `[DEFAULT]` section in `/etc/neutron/metadata_agent.ini` to assist with troubleshooting.

2.



Note

Perform the next two steps on the *controller* node.

3. On the *controller* node, configure Compute to use the metadata service:

Replace `METADATA_SECRET` with the secret you chose for the metadata proxy.

```

# openstack-config --set /etc/nova/nova.conf DEFAULT \
  service_neutron_metadata_proxy true
# openstack-config --set /etc/nova/nova.conf DEFAULT \
  neutron_metadata_proxy_shared_secret METADATA_SECRET

```

4. On the *controller* node, restart the Compute *API* service:

```
# service openstack-nova-api restart
```

To configure the Modular Layer 2 (ML2) plug-in

The ML2 plug-in uses the Open vSwitch (OVS) mechanism (agent) to build virtual networking framework for instances.

- Run the following commands:

Replace `INSTANCE_TUNNELS_INTERFACE_IP_ADDRESS` with the IP address of the instance tunnels network interface on your network node. This guide uses `10.0.1.21` for the IP address of the instance tunnels network interface on the network node.

```

# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 \
  type_drivers gre
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 \
  tenant_network_types gre
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 \
  mechanism_drivers openvswitch
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini
ml2_type_gre \
  tunnel_id_ranges 1:1000
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ovs \
  local_ip INSTANCE_TUNNELS_INTERFACE_IP_ADDRESS

```



```
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ovs \
  tunnel_type gre
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ovs \
  enable_tunneling True
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini
  securitygroup \
  firewall_driver neutron.agent.linux.iptables_firewall.
OVSHybridIptablesFirewallDriver
# openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini
  securitygroup \
  enable_security_group True
```

To configure the Open vSwitch (OVS) service

The OVS service provides the underlying virtual networking framework for instances. The integration bridge `br-int` handles internal instance network traffic within OVS. The external bridge `br-ex` handles external instance network traffic within OVS. The external bridge requires a port on the physical external network interface to provide instances with external network access. In essence, this port bridges the virtual and physical external networks in your environment.

1. Start the OVS service and configure it to start when the system boots:

```
# service openvswitch start
# chkconfig openvswitch on
```

2. Add the integration bridge:

```
# ovs-vsctl add-br br-int
```

3. Add the external bridge:

```
# ovs-vsctl add-br br-ex
```

4. Add a port to the external bridge that connects to the physical external network interface:

Replace `INTERFACE_NAME` with the actual interface name. For example, `eth2` or `ens256`.

```
# ovs-vsctl add-port br-ex INTERFACE_NAME
```



Note

Depending on your network interface driver, you may need to disable *Generic Receive Offload (GRO)* to achieve suitable throughput between your instances and the external network.

To temporarily disable GRO on the external network interface while testing your environment:

```
# ethtool -K INTERFACE_NAME gro off
```

To finalize the installation

1. The Networking service initialization scripts expect a symbolic link `/etc/neutron/plugin.ini` pointing to the configuration file associated with your chosen plug-in.

Using the ML2 plug-in, for example, the symbolic link must point to `/etc/neutron/plugins/ml2/ml2_conf.ini`. If this symbolic link does not exist, create it using the following commands:

```
# ln -s plugins/ml2/ml2_conf.ini /etc/neutron/plugin.ini
```

Due to a packaging bug, the Open vSwitch agent initialization script explicitly looks for the Open vSwitch plug-in configuration file rather than a symbolic link `/etc/neutron/plugin.ini` pointing to the ML2 plug-in configuration file. Run the following commands to resolve this issue:

```
# cp /etc/init.d/neutron-openvswitch-agent /etc/init.d/neutron-  
openvswitch-agent.orig  
# sed -i 's,plugins/openvswitch/ovs_neutron_plugin.ini,plugin.ini,g' /etc/  
init.d/neutron-openvswitch-agent
```

2. Start the Networking services and configure them to start when the system boots:

```
# service neutron-openvswitch-agent start  
# service neutron-l3-agent start  
# service neutron-dhcp-agent start  
# service neutron-metadata-agent start  
# chkconfig neutron-openvswitch-agent on  
# chkconfig neutron-l3-agent on  
# chkconfig neutron-dhcp-agent on  
# chkconfig neutron-metadata-agent on
```

Configure compute node

Prerequisites

Before you configure OpenStack Networking, you must enable certain kernel networking functions.

1. Edit `/etc/sysctl.conf` to contain the following:

```
net.ipv4.conf.all.rp_filter=0  
net.ipv4.conf.default.rp_filter=0
```

2. Implement the changes:

```
# sysctl -p
```

To install the Networking components

- ```
yum install openstack-neutron-ml2 openstack-neutron-openvswitch
```

### To configure the Networking common components

The Networking common component configuration includes the authentication mechanism, message broker, and plug-in.

1. Configure Networking to use the Identity service for authentication:

Replace `NEUTRON_PASS` with the password you chose for the `neutron` user in the Identity service.

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \

```

```
auth_strategy keystone
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 auth_uri http://controller:5000
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 auth_host controller
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 auth_protocol http
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 auth_port 35357
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 admin_tenant_name service
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 admin_user neutron
openstack-config --set /etc/neutron/neutron.conf keystone_auth \
 admin_password NEUTRON_PASS
```

2. Configure Networking to use the message broker:

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
 rpc_backend neutron.openstack.common.rpc.impl_qpid
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
 qpid_hostname controller
```

3. Configure Networking to use the Modular Layer 2 (ML2) plug-in and associated services:

```
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
 core_plugin ml2
openstack-config --set /etc/neutron/neutron.conf DEFAULT \
 service_plugins router
```



### Note

We recommend adding `verbose = True` to the `[DEFAULT]` section in `/etc/neutron/neutron.conf` to assist with troubleshooting.

## To configure the Modular Layer 2 (ML2) plug-in

The ML2 plug-in uses the Open vSwitch (OVS) mechanism (agent) to build the virtual networking framework for instances.

- Run the following commands:

Replace `INSTANCE_TUNNELS_INTERFACE_IP_ADDRESS` with the IP address of the instance tunnels network interface on your compute node. This guide uses `10.0.1.31` for the IP address of the instance tunnels network interface on the first compute node.

```
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 \
 type_drivers gre
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 \
 tenant_network_types gre
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ml2 \
 mechanism_drivers openvswitch
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini
 ml2_type_gre \
 tunnel_id_ranges 1:1000
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ovs \
```

```
local_ip INSTANCE_TUNNELS_INTERFACE_IP_ADDRESS
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ovs \
 tunnel_type gre
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini ovs \
 enable_tunneling True
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini
 securitygroup \
 firewall_driver neutron.agent.linux.iptables_firewall.
OVSHybridIptablesFirewallDriver
openstack-config --set /etc/neutron/plugins/ml2/ml2_conf.ini
 securitygroup \
 enable_security_group True
```

## To configure the Open vSwitch (OVS) service

The OVS service provides the underlying virtual networking framework for instances. The integration bridge `br-int` handles internal instance network traffic within OVS.

1. Start the OVS service and configure it to start when the system boots:

```
service openvswitch start
chkconfig openvswitch on
```

2. Add the integration bridge:

```
ovs-vsctl add-br br-int
```

## To configure Compute to use Networking

By default, most distributions configure Compute to use legacy networking. You must reconfigure Compute to manage networks through Networking.

- Run the following commands:

Replace `NEUTRON_PASS` with the password you chose for the `neutron` user in the Identity service.

```
openstack-config --set /etc/nova/nova.conf DEFAULT \
 network_api_class nova.network.neutronv2.api.API
openstack-config --set /etc/nova/nova.conf DEFAULT \
 neutron_url http://controller:9696
openstack-config --set /etc/nova/nova.conf DEFAULT \
 neutron_auth_strategy keystone
openstack-config --set /etc/nova/nova.conf DEFAULT \
 neutron_admin_tenant_name service
openstack-config --set /etc/nova/nova.conf DEFAULT \
 neutron_admin_username neutron
openstack-config --set /etc/nova/nova.conf DEFAULT \
 neutron_admin_password NEUTRON_PASS
openstack-config --set /etc/nova/nova.conf DEFAULT \
 neutron_admin_auth_url http://controller:35357/v2.0
openstack-config --set /etc/nova/nova.conf DEFAULT \
 linuxnet_interface_driver nova.network.linux_net.LinuxOVSIInterfaceDriver
openstack-config --set /etc/nova/nova.conf DEFAULT \
 firewall_driver nova.virt.firewall.NoopFirewallDriver
openstack-config --set /etc/nova/nova.conf DEFAULT \
 security_group_api neutron
```



## Note

By default, Compute uses an internal firewall service. Since Networking includes a firewall service, you must disable the Compute firewall service by using the `nova.virt.firewall.NoopFirewallDriver` firewall driver.

## To finalize the installation

1. The Networking service initialization scripts expect a symbolic link `/etc/neutron/plugin.ini` pointing to the configuration file associated with your chosen plug-in. Using the ML2 plug-in, for example, the symbolic link must point to `/etc/neutron/plugins/ml2/ml2_conf.ini`. If this symbolic link does not exist, create it using the following commands:

```
ln -s plugins/ml2/ml2_conf.ini /etc/neutron/plugin.ini
```

Due to a packaging bug, the Open vSwitch agent initialization script explicitly looks for the Open vSwitch plug-in configuration file rather than a symbolic link `/etc/neutron/plugin.ini` pointing to the ML2 plug-in configuration file. Run the following commands to resolve this issue:

```
cp /etc/init.d/neutron-openvswitch-agent /etc/init.d/neutron-
openvswitch-agent.orig
sed -i 's,plugins/openvswitch/ovs_neutron_plugin.ini,plugin.ini,g' /etc/
init.d/neutron-openvswitch-agent
```

2. Restart the Compute service:

```
service openstack-nova-compute restart
```

3. Start the Open vSwitch (OVS) agent and configure it to start when the system boots:

```
service neutron-openvswitch-agent start
chkconfig neutron-openvswitch-agent on
```

## Create initial networks

Before launching your first instance, you must create the necessary virtual network infrastructure to which the instance will connect, including the [external network](#) and [tenant network](#). See [Figure 7.1, “Initial networks” \[62\]](#). After creating this infrastructure, we recommend that you [verify connectivity](#) and resolve any issues before proceeding further.



## External network

The external network typically provides internet access for your instances. By default, this network only allows internet access *from* instances using *Network Address Translation (NAT)*. You can enable internet access *to* individual instances using a *floating IP address* and suitable *security group* rules. The `admin` tenant owns this network because it provides external network access for multiple tenants. You must also enable sharing to allow access by those tenants.



### Note

Perform these commands on the controller node.

### To create the external network

1. Source the admin tenant credentials:

```
$ source admin-openrc.sh
```

2. Create the network:

```
$ neutron net-create ext-net --shared --router:external=True
Created a new network:
```

| Field                     | Value                                |
|---------------------------|--------------------------------------|
| admin_state_up            | True                                 |
| id                        | 893aebb9-1c1e-48be-8908-6b947f3237b3 |
| name                      | ext-net                              |
| provider:network_type     | gre                                  |
| provider:physical_network |                                      |
| provider:segmentation_id  | 1                                    |
| router:external           | True                                 |
| shared                    | True                                 |
| status                    | ACTIVE                               |
| subnets                   |                                      |
| tenant_id                 | 54cd044c64d5408b83f843d63624e0d8     |

Like a physical network, a virtual network requires a *subnet* assigned to it. The external network shares the same subnet and *gateway* associated with the physical network connected to the external interface on the network node. You should specify an exclusive slice of this subnet for *router* and floating IP addresses to prevent interference with other devices on the external network.

Replace `FLOATING_IP_START` and `FLOATING_IP_END` with the first and last IP addresses of the range that you want to allocate for floating IP addresses. Replace `EXTERNAL_NETWORK_CIDR` with the subnet associated with the physical network. Replace `EXTERNAL_NETWORK_GATEWAY` with the gateway associated with the physical network, typically the ".1" IP address. You should disable *DHCP* on this subnet because instances do not connect directly to the external network and floating IP addresses require manual assignment.

### To create a subnet on the external network

- Create the subnet:

```
$ neutron subnet-create ext-net --name ext-subnet \
 --allocation-pool start=FLOATING_IP_START,end=FLOATING_IP_END \
 --disable-dhcp --gateway EXTERNAL_NETWORK_GATEWAY EXTERNAL_NETWORK_CIDR
```

For example, using 203.0.113.0/24 with floating IP address range  
203.0.113.101 to 203.0.113.200:

```
$ neutron subnet-create ext-net --name ext-subnet \
 --allocation-pool start=203.0.113.101,end=203.0.113.200 \
 --disable-dhcp --gateway 203.0.113.1 203.0.113.0/24
Created a new subnet:
+-----+
+-----+
Field	Value
+-----+-----+	
allocation_pools	{"start": "203.0.113.101", "end": "203.0.113.200"}
cidr	203.0.113.0/24
dns_nameservers	
enable_dhcp	False
gateway_ip	203.0.113.1
host_routes	
id	9159f0dc-2b63-41cf-bd7a-289309da1391
ip_version	4
ipv6_address_mode	
ipv6_ra_mode	
name	ext-subnet
network_id	893aebb9-1c1e-48be-8908-6b947f3237b3
tenant_id	54cd044c64d5408b83f843d63624e0d8
+-----+
+-----+
```

## Tenant network

The tenant network provides internal network access for instances. The architecture isolates this type of network from other tenants. The `demo` tenant owns this network because it only provides network access for instances within it.



### Note

Perform these commands on the controller node.





```

+-----+
| id | 69d38773-794a-4e49-b887-6de6734e792d |
+-----+
| ip_version | 4 |
+-----+
| ipv6_address_mode | |
+-----+
| ipv6_ra_mode | |
+-----+
| name | demo-subnet |
+-----+
| network_id | ac108952-6096-4243-adf4-bb6615b3de28 |
+-----+
| tenant_id | cdef0071a0194d19ac6bb63802dc9bae |
+-----+

```

A virtual router passes network traffic between two or more virtual networks. Each router requires one or more *interfaces* and/or gateways that provide access to specific networks. In this case, you will create a router and attach your tenant and external networks to it.

### To create a router on the tenant network and attach the external and tenant networks to it

1. Create the router:

```

$ neutron router-create demo-router
Created a new router:
+-----+
| Field | Value |
+-----+
| admin_state_up | True |
| external_gateway_info |
id	635660ae-a254-4feb-8993-295aa9ec6418
name	demo-router
status	ACTIVE
tenant_id	cdef0071a0194d19ac6bb63802dc9bae
+-----+

```

2. Attach the router to the demo tenant subnet:

```

$ neutron router-interface-add demo-router demo-subnet
Added interface bla894fd-ae8-475c-9262-4342afdclb58 to router demo-router.

```

3. Attach the router to the external network by setting it as the gateway:

```

$ neutron router-gateway-set demo-router ext-net
Set gateway for router demo-router

```

### Verify connectivity

We recommend that you verify network connectivity and resolve any issues before proceeding further. Following the external network subnet example using 203.0.113.0/24, the tenant router gateway should occupy the lowest IP address in the floating IP address range, 203.0.113.101. If you configured your external physical network and virtual networks correctly, you should be able to **ping** this IP address from any host on your external physical network.



Replace `INTERFACE_NAME` with the actual interface name for the external network.  
For example, `eth1` or `ens224`.

```
openstack-config --set /etc/nova/nova.conf DEFAULT \
network_api_class nova.network.api.API
openstack-config --set /etc/nova/nova.conf DEFAULT \
security_group_api nova
openstack-config --set /etc/nova/nova.conf DEFAULT \
network_manager nova.network.manager.FlatDHCPManager
openstack-config --set /etc/nova/nova.conf DEFAULT \
firewall_driver nova.virt.libvirt.firewall.IptablesFirewallDriver
openstack-config --set /etc/nova/nova.conf DEFAULT \
network_size 254
openstack-config --set /etc/nova/nova.conf DEFAULT \
allow_same_net_traffic False
openstack-config --set /etc/nova/nova.conf DEFAULT \
multi_host True
openstack-config --set /etc/nova/nova.conf DEFAULT \
send_arp_for_ha True
openstack-config --set /etc/nova/nova.conf DEFAULT \
share_dhcp_address True
openstack-config --set /etc/nova/nova.conf DEFAULT \
force_dhcp_release True
openstack-config --set /etc/nova/nova.conf DEFAULT \
flat_network_bridge br100
openstack-config --set /etc/nova/nova.conf DEFAULT \
flat_interface INTERFACE_NAME
openstack-config --set /etc/nova/nova.conf DEFAULT \
public_interface INTERFACE_NAME
```

2. Start the services and configure them to start when the system boots:

```
service openstack-nova-network start
service openstack-nova-metadata-api start
chkconfig openstack-nova-network on
chkconfig openstack-nova-metadata-api on
```

## Create initial network

Before launching your first instance, you must create the necessary virtual network infrastructure to which the instance will connect. This network typically provides internet access *from* instances. You can enable internet access *to* individual instances using a *floating IP address* and suitable *security group* rules. The *admin* tenant owns this network because it provides external network access for multiple tenants.

This network shares the same *subnet* associated with the physical network connected to the external *interface* on the compute node. You should specify an exclusive slice of this subnet to prevent interference with other devices on the external network.



### Note

Perform these commands on the controller node.

### To create the network

1. Source the *admin* tenant credentials:

```
$ source admin-openrc.sh
```

2. Create the network:

Replace `NETWORK_CIDR` with the subnet associated with the physical network.

```
$ nova network-create demo-net --bridge br100 --multi-host T \
--fixed-range-v4 NETWORK_CIDR
```

For example, using an exclusive slice of `203.0.113.0/24` with IP address range `203.0.113.24` to `203.0.113.32`:

```
$ nova network-create demo-net --bridge br100 --multi-host T \
--fixed-range-v4 203.0.113.24/29
```



### Note

This command provides no output.

3. Verify creation of the network:

```
$ nova net-list
```

| ID                                   | Label    | CIDR            |
|--------------------------------------|----------|-----------------|
| 84b34a65-a762-44d6-8b5e-3b461a53f513 | demo-net | 203.0.113.24/29 |

## Next steps

Your OpenStack environment now includes the core components necessary to launch a basic instance. You can [launch an instance](#) or add more services to your environment in the following chapters.

# 8. Add the dashboard

## Table of Contents

|                             |    |
|-----------------------------|----|
| System requirements .....   | 70 |
| Install the dashboard ..... | 71 |
| Next steps .....            | 72 |

The OpenStack dashboard, also known as [Horizon](#), is a Web interface that enables cloud administrators and users to manage various OpenStack resources and services.

The dashboard enables web-based interactions with the OpenStack Compute cloud controller through the OpenStack APIs.

These instructions show an example deployment configured with an Apache web server.

After you [install and configure the dashboard](#), you can complete the following tasks:

- Customize your dashboard. See section [Customize the dashboard](#) in the *OpenStack Cloud Administrator Guide*.
- Set up session storage. See section [Set up session storage for the dashboard](#) in the *OpenStack Cloud Administrator Guide* for information on user session data.

## System requirements

Before you install the OpenStack dashboard, you must meet the following system requirements:

- OpenStack Compute installation. Enable the Identity Service for user and project management.

Note the URLs of the Identity Service and Compute endpoints.

- Identity Service user with sudo privileges. Because Apache does not serve content from a root user, users must run the dashboard as an Identity Service user with sudo privileges.
- Python 2.6 or 2.7. The Python version must support Django. The Python version should run on any system, including Mac OS X. Installation prerequisites might differ by platform.

Then, install and configure the dashboard on a node that can contact the Identity Service.

Provide users with the following information so that they can access the dashboard through a web browser on their local machine:

- The public IP address from which they can access the dashboard
- The user name and password with which they can access the dashboard

---

Your web browser, and that of your users, must support HTML5 and have cookies and JavaScript enabled.



### Note

To use the VNC client with the dashboard, the browser must support HTML5 Canvas and HTML5 WebSockets.

For details about browsers that support noVNC, see <https://github.com/kanaka/noVNC/blob/master/README.md>, and <https://github.com/kanaka/noVNC/wiki/Browser-support>, respectively.

## Install the dashboard

Before you can install and configure the dashboard, meet the requirements in [the section called “System requirements” \[70\]](#).



### Note

When you install only Object Storage and the Identity Service, even if you install the dashboard, it does not pull up projects and is unusable.

For more information about how to deploy the dashboard, see [deployment topics in the developer documentation](#).

1. Install the dashboard on the node that can contact the Identity Service as root:

```
yum install memcached python-memcached mod_wsgi openstack-dashboard
```

2. Modify the value of `CACHES['default']['LOCATION']` in `/etc/openstack-dashboard/local_settings` to match the ones set in `/etc/sysconfig/memcached`.

Open `/etc/openstack-dashboard/local_settings` and look for this line:

```
CACHES = {
 'default': {
 'BACKEND' : 'django.core.cache.backends.memcached.MemcachedCache',
 'LOCATION' : '127.0.0.1:11211'
 }
}
```



### Notes

- The address and port must match the ones set in `/etc/sysconfig/memcached`.

If you change the memcached settings, you must restart the Apache web server for the changes to take effect.

- You can use options other than memcached option for session storage. Set the session back-end through the `SESSION_ENGINE` option.
- To change the timezone, use the dashboard or edit the `/etc/openstack-dashboard/local_settings` file.





## 9. Add the Block Storage service

### Table of Contents

|                                                    |    |
|----------------------------------------------------|----|
| Block Storage .....                                | 73 |
| Configure a Block Storage service controller ..... | 73 |
| Configure a Block Storage service node .....       | 75 |
| Verify the Block Storage installation .....        | 77 |
| Next steps .....                                   | 78 |

The OpenStack Block Storage service works through the interaction of a series of daemon processes named `cinder-*` that reside persistently on the host machine or machines. You can run the binaries from a single node or across multiple nodes. You can also run them on the same node as other OpenStack services. The following sections introduce Block Storage service components and concepts and show you how to configure and install the Block Storage service.

### Block Storage

The Block Storage service enables management of volumes, volume snapshots, and volume types. It includes the following components:

- `cinder-api`: Accepts API requests and routes them to `cinder-volume` for action.
- `cinder-volume`: Responds to requests to read from and write to the Block Storage database to maintain state, interacting with other processes (like `cinder-scheduler`) through a message queue and directly upon block storage providing hardware or software. It can interact with a variety of storage providers through a driver architecture.
- `cinder-scheduler` daemon: Like the `nova-scheduler`, picks the optimal block storage provider node on which to create the volume.
- Messaging queue: Routes information between the Block Storage service processes.

The Block Storage service interacts with Compute to provide volumes for instances.

### Configure a Block Storage service controller



#### Note

This scenario configures OpenStack Block Storage services on the *Controller node* and assumes that a second node provides storage through the `cinder-volume` service.

For instructions on how to configure the second node, see [the section called "Configure a Block Storage service node" \[75\]](#).

You can configure OpenStack to use various storage systems. This example uses LVM.

1. Install the appropriate packages for the Block Storage service:

```
yum install openstack-cinder
```

2. Configure Block Storage to use your database.

Run the following command to set connection option in the [database] section, which is in the `/etc/cinder/cinder.conf` file, replace `CINDER_DBPASS` with the password for the Block Storage database that you will create in a later step:

```
openstack-config --set /etc/cinder/cinder.conf \
database connection mysql://cinder:CINDER_DBPASS@controller/cinder
```

3. Use the password that you set to log in as root to create a cinder database:

```
mysql -u root -p
mysql> CREATE DATABASE cinder;
mysql> GRANT ALL PRIVILEGES ON cinder.* TO 'cinder'@'localhost' \
IDENTIFIED BY 'CINDER_DBPASS';
mysql> GRANT ALL PRIVILEGES ON cinder.* TO 'cinder'@'%' \
IDENTIFIED BY 'CINDER_DBPASS';
```

4. Create the database tables for the Block Storage service:

```
su -s /bin/sh -c "cinder-manage db sync" cinder
```

5. Create a cinder user.

The Block Storage service uses this user to authenticate with the Identity service.

Use the `service` tenant and give the user the `admin` role:

```
$ keystone user-create --name=cinder --pass=CINDER_PASS --
email=cinder@example.com
$ keystone user-role-add --user=cinder --tenant=service --role=admin
```

6. Edit the `/etc/cinder/cinder.conf` configuration file:

```
openstack-config --set /etc/cinder/cinder.conf DEFAULT \
auth_strategy keystone
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
auth_uri http://controller:5000
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
auth_host controller
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
auth_protocol http
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
auth_port 35357
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
admin_user cinder
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
admin_tenant_name service
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
admin_password CINDER_PASS
```

7. Configure Block Storage to use the Qpid message broker:

```
openstack-config --set /etc/cinder/cinder.conf \
DEFAULT rpc_backend qpid
openstack-config --set /etc/cinder/cinder.conf \
DEFAULT rpc_backend qpid
```

DEFAULT qpid\_hostname controller

- 8. Register the Block Storage service with the Identity service so that other OpenStack services can locate it:

```
$ keystone service-create --name=cinder --type=volume --description="OpenStack Block Storage"
$ keystone endpoint-create \
 --service-id=$(keystone service-list | awk '/ volume / {print $2}') \
 --publicurl=http://controller:8776/v1/%(tenant_id)s \
 --internalurl=http://controller:8776/v1/%(tenant_id)s \
 --adminurl=http://controller:8776/v1/%(tenant_id)s
```

- 9. Register a service and endpoint for version 2 of the Block Storage service API:

```
$ keystone service-create --name=cinderv2 --type=volumev2 --description="OpenStack Block Storage v2"
$ keystone endpoint-create \
 --service-id=$(keystone service-list | awk '/ volumev2 / {print $2}') \
 --publicurl=http://controller:8776/v2/%(tenant_id)s \
 --internalurl=http://controller:8776/v2/%(tenant_id)s \
 --adminurl=http://controller:8776/v2/%(tenant_id)s
```

- 10. Start and configure the Block Storage services to start when the system boots:

```
service openstack-cinder-api start
service openstack-cinder-scheduler start
chkconfig openstack-cinder-api on
chkconfig openstack-cinder-scheduler on
```

## Configure a Block Storage service node

After you configure the services on the controller node, configure a second system to be a Block Storage service node. This node contains the disk that serves volumes.

You can configure OpenStack to use various storage systems. This example uses LVM.

- 1. Use the instructions in [Chapter 2, "Basic environment configuration" \[6\]](#) to configure the system. Note the following differences from the installation instructions for the controller node:
  - Set the host name to `block1` and use `10.0.0.41` as IP address on the management network interface. Ensure that the IP addresses and host names for both controller node and Block Storage service node are listed in the `/etc/hosts` file on each system.
  - Follow the instructions in [the section called "Network Time Protocol \(NTP\)" \[17\]](#) to synchronize from the controller node.
- 2. Create the LVM physical and logical volumes. This guide assumes a second disk `/dev/sdb` that is used for this purpose:

```
pvcreate /dev/sdb
vgcreate cinder-volumes /dev/sdb
```

- 3. Add a filter entry to the `devices` section in the `/etc/lvm/lvm.conf` file to keep LVM from scanning devices used by virtual machines:

```
devices {
 ...
 filter = ["a/sda1/", "a/sdb/", "r/.*/"]
 ...
}
```



### Note

You must add required physical volumes for LVM on the Block Storage host. Run the `pvdisplay` command to get a list or required volumes.

Each item in the filter array starts with either an `a` for accept, or an `r` for reject. The physical volumes that are required on the Block Storage host have names that begin with `a`. The array must end with `"r/.*/"` to reject any device not listed.

In this example, `/dev/sda1` is the volume where the volumes for the operating system for the node reside, while `/dev/sdb` is the volume reserved for `cinder`-volumes.

4. After you configure the operating system, install the appropriate packages for the Block Storage service:

```
yum install openstack-cinder scsi-target-utils
```

5. Copy the `/etc/cinder/cinder.conf` configuration file from the controller, or perform the following steps to set the keystone credentials:

```
openstack-config --set /etc/cinder/cinder.conf DEFAULT \
 auth_strategy keystone
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 auth_uri http://controller:5000
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 auth_host controller
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 auth_protocol http
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 auth_port 35357
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 admin_user cinder
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 admin_tenant_name service
openstack-config --set /etc/cinder/cinder.conf keystone_auth_token \
 admin_password CINDER_PASS
```

6. Configure Block Storage to use the Qpid message broker:

```
openstack-config --set /etc/cinder/cinder.conf \
 DEFAULT rpc_backend qpid
openstack-config --set /etc/cinder/cinder.conf \
 DEFAULT qpid_hostname controller
```

7. Configure Block Storage to use your MySQL database. Edit the `/etc/cinder/cinder.conf` file and add the following key to the `[database]` section. Replace `CINDER_DBPASS` with the password you chose for the Block Storage database:

```
openstack-config --set /etc/cinder/cinder.conf \
 database_connection mysql://cinder:CINDER_DBPASS@controller/cinder
```

8. Configure Block Storage to use the Image Service. Block Storage needs access to images to create bootable volumes. Edit the `/etc/cinder/cinder.conf` file and update the `glance_host` option in the `[DEFAULT]` section:

```
openstack-config --set /etc/cinder/cinder.conf \
 DEFAULT glance_host controller
```

9. Configure the iSCSI target service to discover Block Storage volumes. Add the following line to the beginning of the `/etc/tgt/targets.conf` file, if it is not already present:

```
include /etc/cinder/volumes/*
```

10. Start and configure the Block Storage services to start when the system boots:

```
service openstack-cinder-volume start
service tgt start
chkconfig openstack-cinder-volume on
chkconfig tgt on
```

## Verify the Block Storage installation

To verify that the Block Storage is installed and configured properly, create a new volume.

For more information about how to manage volumes, see the [OpenStack User Guide](#).

1. Source the `demo-openrc.sh` file:

```
$ source demo-openrc.sh
```

2. Use the `cinder create` command to create a new volume:

```
$ cinder create --display-name myVolume 1
```

| Property            | Value                                |
|---------------------|--------------------------------------|
| attachments         | []                                   |
| availability_zone   | nova                                 |
| bootable            | false                                |
| created_at          | 2014-04-17T10:28:19.615050           |
| display_description | None                                 |
| display_name        | myVolume                             |
| encrypted           | False                                |
| id                  | 5e691b7b-12e3-40b6-b714-7f17550db5d1 |
| metadata            | {}                                   |
| size                | 1                                    |
| snapshot_id         | None                                 |
| source_vol_id       | None                                 |
| status              | creating                             |
| volume_type         | None                                 |

3. Make sure that the volume has been correctly created with the `cinder list` command:

```
$ cinder list
```

| ID                                   | Status | Display Name | Size |
|--------------------------------------|--------|--------------|------|
| Volume Type   Bootable   Attached to |        |              |      |

```
+-----+-----+-----+-----+
| 5e691b7b-12e3-40b6-b714-7f17550db5d1 | available | myVolume | 1 |
| None | false | | |
+-----+-----+-----+-----+
```

If the status value is not `available`, the volume creation failed. Check the log files in the `/var/log/cinder/` directory on the controller and volume nodes to get information about the failure.

## Next steps

Your OpenStack environment now includes Block Storage. You can [launch an instance](#) or add more services to your environment in the following chapters.



## System requirements for Object Storage

**Hardware:** OpenStack Object Storage is designed to run on commodity hardware.



### Note

When you install only the Object Storage and Identity Service, you cannot use the dashboard unless you also install Compute and the Image Service.

**Table 10.1. Hardware recommendations**

| Server                                   | Recommended Hardware                                                                                                                           | Notes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Object Storage object servers            | Processor: dual quad core<br>Memory: 8 or 12 GB RAM<br>Disk space: optimized for cost per GB<br>Network: one 1 GB Network Interface Card (NIC) | The amount of disk space depends on how much you can fit into the rack efficiently. You want to optimize these for best cost per GB while still getting industry-standard failure rates. At Rackspace, our storage servers are currently running fairly generic 4U servers with 24 2T SATA drives and 8 cores of processing power. RAID on the storage drives is not required and not recommended. Swift's disk usage pattern is the worst case possible for RAID, and performance degrades very quickly using RAID 5 or 6.<br><br>As an example, Rackspace runs Cloud Files storage servers with 24 2T SATA drives and 8 cores of processing power. Most services support either a worker or concurrency value in the settings. This allows the services to make effective use of the cores available. |
| Object Storage container/account servers | Processor: dual quad core<br>Memory: 8 or 12 GB RAM<br>Network: one 1 GB Network Interface Card (NIC)                                          | Optimized for IOPS due to tracking with SQLite databases.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Object Storage proxy server              | Processor: dual quad core<br>Network: one 1 GB Network Interface Card (NIC)                                                                    | Higher network throughput offers better performance for supporting many API requests.<br><br>Optimize your proxy servers for best CPU performance. The Proxy Services are more CPU and network I/O intensive. If you are using 10 GB networking to the proxy, or are terminating SSL traffic at the proxy, greater CPU power is required.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

**Operating system:** OpenStack Object Storage currently runs on Ubuntu, RHEL, CentOS, Fedora, openSUSE, or SLES.

**Networking:** 1 Gbps or 10 Gbps is suggested internally. For OpenStack Object Storage, an external network should connect the outside world to the proxy servers, and the storage network is intended to be isolated on a private network or multiple private networks.

**Database:** For OpenStack Object Storage, a SQLite database is part of the OpenStack Object Storage container and account management process.

**Permissions:** You can install OpenStack Object Storage either as root or as a user with sudo permissions if you configure the sudoers file to enable all the permissions.

## Plan networking for Object Storage

For both conserving network resources and ensuring that network administrators understand the needs for networks and public IP addresses for providing access to the








**OpenStack Object Storage**

Stores container databases, account databases, and stored objects



## Install Object Storage

Though you can install OpenStack Object Storage for development or testing purposes on one server, a multiple-server installation enables the high availability and redundancy you want in a production distributed object storage system.

To perform a single-node installation for development purposes from source code, use the Swift All In One instructions (Ubuntu) or DevStack (multiple distros). See [http://swift.openstack.org/development\\_saio.html](http://swift.openstack.org/development_saio.html) for manual instructions or <http://devstack.org> for all-in-one including authentication with the Identity Service (keystone) v2.0 API.



### Warning

In this guide we recommend installing and configuring the Identity service so that it implements Identity API v2.0. The Object Storage service is unaware of domains when implementing Access Control Lists (ACLs), so you must use the v2.0 API to avoid having identical user names in different domains, which would enable two users to access the same objects.

## Before you begin

Have a copy of the operating system installation media available if you are installing on a new server.

These steps assume you have set up repositories for packages for your operating system as shown in [OpenStack Packages](#).

This document demonstrates how to install a cluster by using the following types of nodes:

- One proxy node which runs the swift-proxy-server processes. The proxy server proxies requests to the appropriate storage nodes.





```
[account]
max connections = 2
path = /srv/node/
read only = false
lock file = /var/lock/account.lock

[container]
max connections = 2
path = /srv/node/
read only = false
lock file = /var/lock/container.lock

[object]
max connections = 2
path = /srv/node/
read only = false
lock file = /var/lock/object.lock
```

- (Optional) If you want to separate rsync and replication traffic to replication network, set `STORAGE_REPLICATION_NET_IP` instead of `STORAGE_LOCAL_NET_IP`:

```
address = STORAGE_REPLICATION_NET_IP
```

- Edit the following line in `/etc/xinetd.d/rsync`:

```
disable = no
```

- Start the `xinetd` service:

```
service xinetd start
```



### Note

The rsync service requires no authentication, so run it on a local, private network.

- Create the swift recon cache directory and set its permissions:

```
mkdir -p /var/swift/recon
chown -R swift:swift /var/swift/recon
```

## Install and configure the proxy node

The proxy server takes each request and looks up locations for the account, container, or object and routes the requests correctly. The proxy server also handles API requests. You enable account management by configuring it in the `/etc/swift/proxy-server.conf` file.



### Note

The Object Storage processes run under a separate user and group, set by configuration options, and referred to as `swift:swift`. The default user is `swift`.

- Install `swift-proxy` service:

```
yum install openstack-swift-proxy memcached python-swiftclient python-keystone-auth-token
```











# 11. Add the Orchestration service

## Table of Contents

|                                                     |    |
|-----------------------------------------------------|----|
| Orchestration service overview .....                | 91 |
| Install the Orchestration service .....             | 91 |
| Verify the Orchestration service installation ..... | 93 |
| Next steps .....                                    | 95 |

Use the Orchestration module to create cloud resources using a template language called HOT. The integrated project name is Heat.

## Orchestration service overview

The Orchestration service provides a template-based orchestration for describing a cloud application by running OpenStack API calls to generate running cloud applications. The software integrates other core components of OpenStack into a one-file template system. The templates enable you to create most OpenStack resource types, such as instances, floating IPs, volumes, security groups, users, and so on. Also, provides some more advanced functionality, such as instance high availability, instance auto-scaling, and nested stacks. By providing very tight integration with other OpenStack core projects, all OpenStack core projects could receive a larger user base.

The service enables deployers to integrate with the Orchestration service directly or through custom plug-ins.

The Orchestration service consists of the following components:

- `heat` command-line client. A CLI that communicates with the `heat-api` to run AWS CloudFormation APIs. End developers could also use the Orchestration REST API directly.
- `heat-api` component. Provides an OpenStack-native REST API that processes API requests by sending them to the `heat-engine` over RPC.
- `heat-api-cfn` component. Provides an AWS Query API that is compatible with AWS CloudFormation and processes API requests by sending them to the `heat-engine` over RPC.
- `heat-engine`. Orchestrates the launching of templates and provides events back to the API consumer.

## Install the Orchestration service

1. Install the Orchestration module on the controller node:

```
yum install openstack-heat-api openstack-heat-engine \
openstack-heat-api-cfn
```

2. In the configuration file, specify the location of the database where the Orchestration service stores data. These examples use a MySQL database with a `heat` user on the controller node. Replace `HEAT_DBPASS` with the password for the database user:

```
openstack-config --set /etc/heat/heat.conf \
database connection mysql://heat:HEAT_DBPASS@controller/heat
```

3. Use the password that you set previously to log in as `root` and create a heat database user:

```
$ mysql -u root -p
mysql> CREATE DATABASE heat;
mysql> GRANT ALL PRIVILEGES ON heat.* TO 'heat'@'localhost' \
IDENTIFIED BY 'HEAT_DBPASS';
mysql> GRANT ALL PRIVILEGES ON heat.* TO 'heat'@'% ' \
IDENTIFIED BY 'HEAT_DBPASS';
```

4. Create the heat service tables:

```
su -s /bin/sh -c "heat-manage db_sync" heat
```



### Note

Ignore DeprecationWarning errors.

5. Configure the Orchestration Service to use the Qpid message broker:

```
openstack-config --set /etc/heat/heat.conf DEFAULT
qpid_hostname controller
```

6. Create a `heat` user that the Orchestration service can use to authenticate with the Identity Service. Use the `service` tenant and give the user the `admin` role:

```
$ keystone user-create --name=heat --pass=HEAT_PASS \
--email=heat@example.com
$ keystone user-role-add --user=heat --tenant=service --role=admin
```

7. Run the following commands to configure the Orchestration service to authenticate with the Identity service:

```
openstack-config --set /etc/heat/heat.conf keystone_authtoken \
auth_uri http://controller:5000/v2.0
openstack-config --set /etc/heat/heat.conf keystone_authtoken \
auth_port 35357
openstack-config --set /etc/heat/heat.conf keystone_authtoken \
auth_protocol http
openstack-config --set /etc/heat/heat.conf keystone_authtoken \
admin_tenant_name service
openstack-config --set /etc/heat/heat.conf keystone_authtoken \
admin_user heat
openstack-config --set /etc/heat/heat.conf keystone_authtoken \
admin_passwordHEAT_PASS
openstack-config --set /etc/heat/heat.conf ec2authtoken \
auth_uri http://controller:5000/v2.0
```

8. Register the Heat and CloudFormation APIs with the Identity Service so that other OpenStack services can locate these APIs. Register the services and specify the endpoints:

```
$ keystone service-create --name=heat --type=orchestration \
--description="Orchestration"
$ keystone endpoint-create \
--service-id=$(keystone service-list | grep ^heat | cut -d'|' -f1) \
--region=$(keystone endpoint-list | grep ^heat | cut -d'|' -f1) \
--url=http://controller:5000/v2.0
```







## 12. Add the Telemetry module

### Table of Contents

|                                                          |     |
|----------------------------------------------------------|-----|
| Telemetry .....                                          | 96  |
| Install the Telemetry module .....                       | 97  |
| Install the Compute agent for Telemetry .....            | 99  |
| Configure the Image Service for Telemetry .....          | 100 |
| Add the Block Storage service agent for Telemetry .....  | 101 |
| Configure the Object Storage service for Telemetry ..... | 101 |
| Verify the Telemetry installation .....                  | 102 |
| Next steps .....                                         | 103 |

Telemetry provides a framework for monitoring and metering the OpenStack cloud. It is also known as the Ceilometer project.

### Telemetry

The Telemetry module:

- Efficiently collects the metering data about the CPU and network costs.
- Collects data by monitoring notifications sent from services or by polling the infrastructure.
- Configures the type of collected data to meet various operating requirements. Accessing and inserting the metering data through the REST API.
- Expands the framework to collect custom usage data by additional plug-ins.
- Produces signed metering messages that cannot be repudiated.

The system consists of the following basic components:

- A compute agent (`ceilometer-agent-compute`). Runs on each compute node and polls for resource utilization statistics. There may be other types of agents in the future, but for now we will focus on creating the compute agent.
- A central agent (`ceilometer-agent-central`). Runs on a central management server to poll for resource utilization statistics for resources not tied to instances or compute nodes.
- A collector (`ceilometer-collector`). Runs on one or more central management servers to monitor the message queues (for notifications and for metering data coming from the agent). Notification messages are processed and turned into metering messages and sent back out onto the message bus using the appropriate topic. Telemetry messages are written to the data store without modification.
- An alarm notifier (`ceilometer-alarm-notifier`). Runs on one or more central management servers to allow setting alarms based on threshold evaluation for a collection of samples.





5. Create the database and a `ceilometer` database user:

```
mongo --host controller --eval '
db = db.getSiblingDB("ceilometer");
db.addUser({user: "ceilometer",
 pwd: "CEILOMETER_DBPASS",
 roles: ["readWrite", "dbAdmin"]})'
```

6. Configure the Telemetry service to use the database:

```
openstack-config --set /etc/ceilometer/ceilometer.conf \
database connection mongoddb://
ceilometer:CEILOMETER_DBPASS@controller:27017/ceilometer
```

7. You must define a secret key that is used as a shared secret among Telemetry service nodes. Use `openssl` to generate a random token and store it in the configuration file:

```
CEILOMETER_TOKEN=$(openssl rand -hex 10)
echo $CEILOMETER_TOKEN
openstack-config --set /etc/ceilometer/ceilometer.conf publisher
metering_secret $CEILOMETER_TOKEN
```

8. Configure the Qpid access:

```
openstack-config --set /etc/ceilometer/ceilometer.conf \
DEFAULT rpc_backend ceilometer.openstack.common.rpc.impl_qpid
```

9. Create a `ceilometer` user that the Telemetry service uses to authenticate with the Identity Service. Use the `service` tenant and give the user the `admin` role:

```
$ keystone user-create --name=ceilometer --pass=CEILOMETER_PASS --
email=ceilometer@example.com
$ keystone user-role-add --user=ceilometer --tenant=service --role=admin
```

10. Configure the Telemetry service to authenticate with the Identity Service.

Set the `auth_strategy` value to `keystone` in the `/etc/ceilometer/ceilometer.conf` file:

```
openstack-config --set /etc/ceilometer/ceilometer.conf \
DEFAULT auth_strategy keystone
```

11. Add the credentials to the configuration files for the Telemetry service:

```
openstack-config --set /etc/ceilometer/ceilometer.conf \
keystone_authtoken auth_host controller
openstack-config --set /etc/ceilometer/ceilometer.conf \
keystone_authtoken admin_user ceilometer
openstack-config --set /etc/ceilometer/ceilometer.conf \
keystone_authtoken admin_tenant_name service
openstack-config --set /etc/ceilometer/ceilometer.conf \
keystone_authtoken auth_protocol http
openstack-config --set /etc/ceilometer/ceilometer.conf \
keystone_authtoken auth_uri http://controller:5000
openstack-config --set /etc/ceilometer/ceilometer.conf \
keystone_authtoken admin_password CEILOMETER_PASS
openstack-config --set /etc/ceilometer/ceilometer.conf \
service_credentials os_auth_url http://controller:5000/v2.0
openstack-config --set /etc/ceilometer/ceilometer.conf \
service_credentials os_username ceilometer
```





```
openstack-config --set /etc/glance/glance-api.conf DEFAULT
notification_driver messaging
openstack-config --set /etc/glance/glance-api.conf DEFAULT rpc_backend
qpuid
```

2. Restart the Image Services with their new settings:

```
service openstack-glance-api restart
service openstack-glance-registry restart
```

## Add the Block Storage service agent for Telemetry

1. To retrieve volume samples, you must configure the Block Storage service to send notifications to the bus.

Run the following commands on the controller and volume nodes:

```
openstack-config --set /etc/cinder/cinder.conf DEFAULT control_exchange
cinder
openstack-config --set /etc/cinder/cinder.conf DEFAULT
notification_driver cinder.openstack.common.notifier.rpc_notifier
```

2. Restart the Block Storage services with their new settings.

On the controller node:

```
service openstack-cinder-api restart
service openstack-cinder-scheduler restart
```

On the volume node:

```
service openstack-cinder-volume restart
```

## Configure the Object Storage service for Telemetry

1. To retrieve object store statistics, the Telemetry service needs access to Object Storage with the ResellerAdmin role. Give this role to your `os_username` user for the `os_tenant_name` tenant:

```
$ keystone role-create --name=ResellerAdmin
+-----+-----+
| Property | Value |
+-----+-----+
| id | 462fa46c13fd4798a95a3bfbe27b5e54 |
| name | ResellerAdmin |
+-----+-----+

$ keystone user-role-add --tenant service --user ceilometer \
 --role 462fa46c13fd4798a95a3bfbe27b5e54
```

2. You must also add the Telemetry middleware to Object Storage to handle incoming and outgoing traffic. Add these lines to the `/etc/swift/proxy-server.conf` file:

```
[filter:ceilometer]
use = egg:ceilometer#swift
```



```
$ ceilometer statistics -m image.download -p 60
```

```
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Period | Period Start | Period End | Count | Min
| Max | Sum | Avg | Duration | Duration Start
| | Duration End | | |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 60 | 2013-11-18T18:08:50 | 2013-11-18T18:09:50 | 1 | 13167616.0
| 13167616.0 | 13167616.0 | 13167616.0 | 0.0 | 2013-11-18T18:09:05.
334000 | 2013-11-18T18:09:05.334000 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```

## Next steps

Your OpenStack environment now includes Telemetry. You can [launch an instance](#) or add more services to your environment in the previous chapters.

## 13. Add the Database service

### Table of Contents

|                                                |     |
|------------------------------------------------|-----|
| Database service overview .....                | 104 |
| Install the Database service .....             | 105 |
| Verify the Database service installation ..... | 108 |

Use the *Database module* to create cloud database resources. The integrated project name is *trove*.



#### Warning

This chapter is a work in progress. It may contain incorrect information, and will be updated frequently.

### Database service overview

The Database service provides scalable and reliable cloud provisioning functionality for both relational and non-relational database engines. Users can quickly and easily utilize database features without the burden of handling complex administrative tasks. Cloud users and database administrators can provision and manage multiple database instances as needed.

The Database service provides resource isolation at high performance levels, and automates complex administrative tasks such as deployment, configuration, patching, backups, restores, and monitoring.

**Process flow example.** Here is a high-level process flow example for using Database services:

1. Administrator sets up infrastructure:
  - a. OpenStack administrator installs the Database service.
  - b. She creates one image for each type of database the administrator wants to have (one for MySQL, one for MongoDB, and so on).
  - c. OpenStack administrator updates the datastore to use the new images, using the **trove-manage** command.
2. End user uses database service:
  - a. Now that the basic infrastructure is set up, an end user can create a Trove instance (database) whenever the user wants, using the **trove create** command.
  - b. The end user gets the IP address of the Trove instance by using the **trove list** command to get the ID of the instance, and then using the **trove show instanceID** command to get the IP address.
  - c. The end user can now access the Trove instance using typical database access commands. MySQL example:



```
$ mysql -u myuser -pmypass -h trove_ip_address mydb
```

**Components:** The Database service includes the following components:

- `python-troveclient` command-line client. A CLI that communicates with the `trove-api` component.
- `trove-api` component. Provides an OpenStack-native RESTful API that supports JSON to provision and manage Trove instances.
- `trove-conductor` service. Runs on the host, and receives messages from guest instances that want to update information on the host.
- `trove-taskmanager` service. Instruments the complex system flows that support provisioning instances, managing the lifecycle of instances, and performing operations on instances.
- `trove-guestagent` service. Runs within the guest instance. Manages and performs operations on the database itself.

## Install the Database service

This procedure installs the Database module on the controller node.

**Prerequisites.** This chapter assumes that you already have a working OpenStack environment with at least the following components installed: Compute, Image Service, Identity.

To install the Database module on the controller:

1. Install required packages:

```
yum install openstack-trove python-troveclient
```

2. Prepare OpenStack:

- a. Source the `admin-openrc.sh` file.

```
$ source ~/admin-openrc.sh
```

- b. Create a `trove` user that Compute uses to authenticate with the Identity service. Use the `service` tenant and give the user the `admin` role:

```
$ keystone user-create --name=trove --pass=TROVE_PASS \
--email=trove@example.com
$ keystone user-role-add --user=trove --tenant=service --role=admin
```

3. Edit the following configuration files, taking the below actions for each file:

- `trove.conf`
- `trove-taskmanager.conf`
- `trove-conductor.conf`



```
[DEFAULT]
....
Configuration options for talking to nova via the novaclient.
These options are for an admin user in your keystone config.
It proxy's the token received from the user to send to nova via this
admin users creds,
basically acting like the client via that proxy token.
nova_proxy_admin_user = admin
nova_proxy_admin_pass = ADMIN_PASS
nova_proxy_admin_tenant_name = service
...
```

7. Prepare the trove admin database:

```
$ mysql -u root -p
mysql> CREATE DATABASE trove;
mysql> GRANT ALL PRIVILEGES ON trove.* TO trove@'localhost' IDENTIFIED BY
'TROVE_DBPASS';
mysql> GRANT ALL PRIVILEGES ON trove.* TO trove@'%' IDENTIFIED BY
'TROVE_DBPASS';
```

8. Prepare the Database service:

- a. Initialize the database:

```
su -s /bin/sh -c "trove-manage db_sync" trove
```

- b. Create a datastore. You need to create a separate datastore for each type of database you want to use, for example, MySQL, MongoDB, Cassandra. This example shows you how to create a datastore for a MySQL database:

```
su -s /bin/sh -c "trove-manage datastore_update mysql ''" trove
```

9. Create a trove image.

Create an image for the type of database you want to use, for example, MySQL, MongoDB, Cassandra.

This image must have the trove guest agent installed, and it must have the `trove-guestagent.conf` file configured to connect to your OpenStack environment. To correctly configure the `trove-guestagent.conf` file, follow these steps on the guest instance you are using to build your image:

- Add the following lines to `trove-guestagent.conf`:

```
rpc_backend = trove.openstack.common.rpc.impl_qpid
qpid_host = controller
nova_proxy_admin_user = admin
nova_proxy_admin_pass = ADMIN_PASS
nova_proxy_admin_tenant_name = service
trove_auth_url = http://controller:35357/v2.0
```

10. Update the datastore to use the new image, using the `trove-manage` command.

This example shows you how to create a MySQL 5.5 datastore:

```
trove-manage --config-file=/etc/trove/trove.conf
datastore_version_update \
```



# 14. Launch an instance

## Table of Contents

|                                                                |     |
|----------------------------------------------------------------|-----|
| Launch an instance with OpenStack Networking (neutron) .....   | 109 |
| Launch an instance with legacy networking (nova-network) ..... | 115 |

An instance is a VM that OpenStack provisions on a compute node. This guide shows you how to launch a minimal instance using the *CirrOS* image that you added to your environment in the [Chapter 5, “Configure the Image Service” \[35\]](#) chapter. In these steps, you use the command-line interface (CLI) on your controller node or any system with the appropriate OpenStack client libraries. To use the dashboard, see the [OpenStack User Guide](#).

Launch an instance using [OpenStack Networking \(neutron\)](#) or [legacy networking \(nova-network\)](#). For more information, see the [OpenStack User Guide](#).



### Note

These steps reference example components created in previous chapters. You must adjust certain values such as IP addresses to match your environment.

## Launch an instance with OpenStack Networking (neutron)

### To generate a keypair

Most cloud images support *public key authentication* rather than conventional username/password authentication. Before launching an instance, you must generate a public/private key pair using `ssh-keygen` and add the public key to your OpenStack environment.

1. Source the `demo` tenant credentials:

```
$ source demo-openrc.sh
```

2. Generate a key pair:

```
$ ssh-keygen
```

3. Add the public key to your OpenStack environment:

```
$ nova keypair-add --pub-key ~/.ssh/id_rsa.pub demo-key
```



### Note

This command provides no output.

4. Verify addition of the public key:

```
$ nova keypair-list
```

```
+-----+-----+-----+-----+-----+
```

| Name     | Fingerprint                                     |
|----------|-------------------------------------------------|
| demo-key | 6c:74:ec:3a:08:05:4e:9e:21:22:a6:dd:b2:62:b8:28 |

### To launch an instance

To launch an instance, you must at least specify the flavor, image name, network, security group, key, and instance name.

1. A flavor specifies a virtual resource allocation profile which includes processor, memory, and storage.

List available flavors:

```
$ nova flavor-list
+-----+-----+-----+-----+-----+-----+-----+
| ID | Name | Memory_MB | Disk | Ephemeral | Swap | VCPUs |
| RXTX_Factor | Is_Public |
+-----+-----+-----+-----+-----+-----+-----+
| 1 | ml.tiny | 512 | 1 | 0 | | 1 | 1.0
| | True |
| 2 | ml.small | 2048 | 20 | 0 | | 1 | 1.0
| | True |
| 3 | ml.medium | 4096 | 40 | 0 | | 2 | 1.0
| | True |
| 4 | ml.large | 8192 | 80 | 0 | | 4 | 1.0
| | True |
| 5 | ml.xlarge | 16384 | 160 | 0 | | 8 | 1.0
| | True |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
```

Your first instance uses the `ml.tiny` flavor.



### Note

You can also reference a flavor by ID.

2. List available images:

```
$ nova image-list
+-----+-----+-----+-----+-----+-----+-----+
| ID | Name | Status |
| Server |
+-----+-----+-----+-----+-----+-----+-----+
| acafc7c0-40aa-4026-9673-b879898e1fc2 | cirros-0.3.2-x86_64 | ACTIVE |
| |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
```

Your first instance uses the `cirros-0.3.2-x86_64` image.

3. List available networks:















```
+-----+
+-----+
| acafc7c0-40aa-4026-9673-b879898e1fc2 | cirros-0.3.2-x86_64 | ACTIVE |
| | | |
+-----+
+-----+
```

Your first instance uses the `cirros-0.3.2-x86_64` image.

3. List available networks:



### Note

You must source the `admin` tenant credentials for this step and then source the `demo` tenant credentials for the remaining steps.

```
$ source admin-openrc.sh
```

```
$ nova net-list
+-----+
| ID | Label | CIDR |
+-----+
| 7f849be3-4494-495a-95a1-0f99ccb884c4 | demo-net | 203.0.113.24/29 |
+-----+
```

Your first instance uses the `demo-net` tenant network. However, you must reference this network using the ID instead of the name.

4. List available security groups:

```
$ nova secgroup-list
+-----+
| Id | Name | Description |
+-----+
| ad8d4ea5-3cad-4f7d-b164-ada67ec59473 | default | default |
+-----+
```

Your first instance uses the `default` security group. By default, this security group implements a firewall that blocks remote access to instances. If you would like to permit remote access to your instance, launch it and then [configure remote access](#).

5. Launch the instance:

Replace `DEMO_NET_ID` with the ID of the `demo-net` tenant network.

```
$ nova boot --flavor m1.tiny --image cirros-0.3.2-x86_64 --nic net-
id=DEMO_NET_ID \
 --security-group default --key-name demo-key demo-instance1
+-----+
| Property | Value |
+-----+
| OS-DCF:diskConfig | MANUAL |
| OS-EXT-AZ:availability_zone | nova |
+-----+
```

```

| OS-EXT-STS:power_state | 0
| OS-EXT-STS:task_state | scheduling
| OS-EXT-STS:vm_state | building
| OS-SRV-USG:launched_at | -
| OS-SRV-USG:terminated_at | -
| accessIPv4 |
| accessIPv6 |
| adminPass | ThZqrg7ach78
| config_drive |
| created | 2014-04-10T00:09:16Z
| flavor | ml.tiny (1)
| hostId |
| id | 45ea195c-
c469-43eb-83db-1a663bbad2fc
| image | cirros-0.3.2-x86_64
(acafc7c0-40aa-4026-9673-b879898e1fc2)
| key_name | demo-key
| metadata | {}
| name | demo-instance1
| os-extended-volumes:volumes_attached | []
| progress | 0
| security_groups | default
| status | BUILD
| tenant_id | 93849608fe3d462ca9fa0e5dbfd4d040
| updated | 2014-04-10T00:09:16Z
| user_id | 8397567baf4746cca7a1e608677c3b23
+-----+
+-----+

```

6. Check the status of your instance:

```

$ nova list
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| ID | Name | Status | Task
State | Power State | Networks |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+

```



| IP Protocol | From Port | To Port | IP Range  | Source Group |
|-------------|-----------|---------|-----------|--------------|
| icmp        | -1        | -1      | 0.0.0.0/0 |              |

- b. Permit secure shell (SSH) access:

```
$ nova secgroup-add-rule default tcp 22 22 0.0.0.0/0
```

| IP Protocol | From Port | To Port | IP Range  | Source Group |
|-------------|-----------|---------|-----------|--------------|
| tcp         | 22        | 22      | 0.0.0.0/0 |              |

2. Verify network connectivity using **ping** from the controller node or any host on the external network:

```
$ ping -c 4 203.0.113.26
PING 203.0.113.26 (203.0.113.26) 56(84) bytes of data.
64 bytes from 203.0.113.26: icmp_req=1 ttl=63 time=3.18 ms
64 bytes from 203.0.113.26: icmp_req=2 ttl=63 time=0.981 ms
64 bytes from 203.0.113.26: icmp_req=3 ttl=63 time=1.06 ms
64 bytes from 203.0.113.26: icmp_req=4 ttl=63 time=0.929 ms

--- 203.0.113.26 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3002ms
rtt min/avg/max/mdev = 0.929/1.539/3.183/0.951 ms
```

3. Access your instance using SSH from the controller node or any host on the external network:

```
$ ssh cirros@203.0.113.26
The authenticity of host '203.0.113.26 (203.0.113.26)' can't be
established.
RSA key fingerprint is ed:05:e9:e7:52:a0:ff:83:68:94:c7:d1:f2:f8:e2:e9.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '203.0.113.26' (RSA) to the list of known
hosts.
$
```



### Note

If your host does not contain the public/private key pair created in an earlier step, SSH prompts for the default password associated with the `cirros` user.

If your instance does not launch or seem to work as you expect, see the [OpenStack Operations Guide](#) for more information or use one of the [many other options](#) to seek assistance. We want your environment to work!





# Appendix B. Community support

## Table of Contents

|                                       |     |
|---------------------------------------|-----|
| Documentation .....                   | 122 |
| ask.openstack.org .....               | 123 |
| OpenStack mailing lists .....         | 123 |
| The OpenStack wiki .....              | 124 |
| The Launchpad Bugs area .....         | 124 |
| The OpenStack IRC channel .....       | 125 |
| Documentation feedback .....          | 125 |
| OpenStack distribution packages ..... | 125 |

The following resources are available to help you run and use OpenStack. The OpenStack community constantly improves and adds to the main features of OpenStack, but if you have any questions, do not hesitate to ask. Use the following resources to get OpenStack support, and troubleshoot your installations.

## Documentation

For the available OpenStack documentation, see [docs.openstack.org](http://docs.openstack.org).

To provide feedback on documentation, join and use the <openstack-docs@lists.openstack.org> mailing list at [OpenStack Documentation Mailing List](#), or [report a bug](#).

The following books explain how to install an OpenStack cloud and its associated components:

- [Installation Guide for Debian 7.0](#)
- [Installation Guide for openSUSE and SUSE Linux Enterprise Server](#)
- [Installation Guide for Red Hat Enterprise Linux, CentOS, and Fedora](#)
- [Installation Guide for Ubuntu 12.04/14.04 \(LTS\)](#)

The following books explain how to configure and run an OpenStack cloud:

- [Cloud Administrator Guide](#)
- [Configuration Reference](#)
- [Operations Guide](#)
- [High Availability Guide](#)
- [Security Guide](#)



---

## The OpenStack wiki

The [OpenStack wiki](#) contains a broad range of topics but some of the information can be difficult to find or is a few pages deep. Fortunately, the wiki search feature enables you to search by title or content. If you search for specific information, such as about networking or nova, you can find a large amount of relevant material. More is being added all the time, so be sure to check back often. You can find the search box in the upper-right corner of any OpenStack wiki page.

## The Launchpad Bugs area

The OpenStack community values your set up and testing efforts and wants your feedback. To log a bug, you must sign up for a Launchpad account at <https://launchpad.net/+login>. You can view existing bugs and report bugs in the Launchpad Bugs area. Use the search feature to determine whether the bug has already been reported or already been fixed. If it still seems like your bug is unreported, fill out a bug report.

Some tips:

- Give a clear, concise summary.
- Provide as much detail as possible in the description. Paste in your command output or stack traces, links to screen shots, and any other information which might be useful.
- Be sure to include the software and package versions that you are using, especially if you are using a development branch, such as, "Juno release" vs `git commit bc79c3ecc55929bac585d04a03475b72e06a3208`.
- Any deployment-specific information is helpful, such as whether you are using Ubuntu 14.04 or are performing a multi-node installation.

The following Launchpad Bugs areas are available:

- [Bugs: OpenStack Block Storage \(cinder\)](#)
- [Bugs: OpenStack Compute \(nova\)](#)
- [Bugs: OpenStack Dashboard \(horizon\)](#)
- [Bugs: OpenStack Identity \(keystone\)](#)
- [Bugs: OpenStack Image Service \(glance\)](#)
- [Bugs: OpenStack Networking \(neutron\)](#)
- [Bugs: OpenStack Object Storage \(swift\)](#)
- [Bugs: Bare Metal \(ironic\)](#)
- [Bugs: Data Processing Service \(sahara\)](#)
- [Bugs: Database Service \(trove\)](#)



---

# Glossary

## API

Application programming interface.

## API endpoint

The daemon, worker, or service that a client communicates with to access an API. API endpoints can provide any number of services, such as authentication, sales data, performance metrics, Compute VM commands, census data, and so on.

## Block Storage

The OpenStack core project that enables management of volumes, volume snapshots, and volume types. The project name of Block Storage is cinder.

## CirrOS

A minimal Linux distribution designed for use as a test image on clouds such as OpenStack.

## cloud controller node

A node that runs network, volume, API, scheduler, and image services. Each service may be broken out into separate nodes for scalability or availability.

## Compute

The OpenStack core project that provides compute services. The project name of the Compute service is nova.

## compute node

A node that runs the `nova-compute` daemon which manages VM instances that provide a wide range of services such as a web applications and analytics.

## controller node

Alternative term for a cloud controller node.

## Database Service

An integrated project that provide scalable and reliable Cloud Database-as-a-Service functionality for both relational and non-relational database engines. The project name of Database Service is trove.

## DHCP

Dynamic Host Configuration Protocol. A network protocol that configures devices that are connected to a network so that they can communicate on that network by using the Internet Protocol (IP). The protocol is implemented in a client-server model where DHCP clients request configuration data such as, an IP address, a default route, and one or more DNS server addresses from a DHCP server.

## DHCP agent

OpenStack Networking agent that provides DHCP services for virtual networks.

## endpoint

See API endpoint.

## external network

A network segment typically used for instance Internet access.

Used to restrict communications between hosts and/or nodes, implemented in Compute using iptables, arptables, ip6tables, and etables.

#### flat network

The Network Controller provides virtual networks to enable compute servers to interact with each other and with the public network. All machines must have a public and private network interface. A flat network is a private network interface, which is controlled by the flat\_interface option with flat managers.

#### floating IP address

An IP address that a project can associate with a VM so that the instance has the same public IP address each time that it boots. You create a pool of floating IP addresses and assign them to instances as they are launched to maintain a consistent IP address for maintaining DNS assignment.

#### gateway

An IP address, typically assigned to a router, that passes network traffic between different networks.

#### Generic Receive Offload (GRO)

Feature of certain network interface drivers that combines many smaller received packets into a large packet before delivery to the kernel IP stack.

#### generic routing encapsulation (GRE)

Protocol that encapsulates a wide variety of network layer protocols inside virtual point-to-point links.

#### hypervisor

Software that arbitrates and controls VM access to the actual underlying hardware.

#### IaaS

Infrastructure-as-a-Service. IaaS is a provisioning model in which an organization outsources physical components of a data center such as storage, hardware, servers and networking components. A service provider owns the equipment and is responsible for housing, operating and maintaining it. The client typically pays on a per-use basis. IaaS is a model for providing cloud services.

#### ICMP

Internet Control Message Protocol, used by network devices for control messages. For example, **ping** uses ICMP to test connectivity.

#### Identity Service

The OpenStack core project that provides a central directory of users mapped to the OpenStack services they can access. It also registers endpoints for OpenStack services. It acts as a common authentication system. The project name of the Identity Service is keystone.

#### Internet protocol (IP)

Principal communications protocol in the internet protocol suite for relaying datagrams across network boundaries.

#### Image Service

An OpenStack core project that provides discovery, registration, and delivery services for disk and server images. The project name of the Image Service is glance.

---

**instance tunnels network**

A network segment used for instance traffic tunnels between compute nodes and the network node.

**interface**

A physical or virtual device that provides connectivity to another device or medium.

**jumbo frames**

Feature in modern Ethernet networks that supports frames up to approximately 9000 bytes.

**kernel-based VM (KVM)**

An OpenStack-supported hypervisor.

**Layer-3 (L3) agent**

OpenStack Networking agent that provides layer-3 (routing) services for virtual networks.

**load balancer**

A load balancer is a logical device that belongs to a cloud account. It is used to distribute workloads between multiple back-end systems or services, based on the criteria defined as part of its configuration.

**Logical Volume Manager (LVM)**

Provides a method of allocating space on mass-storage devices that is more flexible than conventional partitioning schemes.

**management network**

A network segment used for administration, not accessible to the public Internet.

**maximum transmission unit (MTU)**

Maximum frame or packet size for a particular network medium. Typically 1500 bytes for Ethernet networks.

**message broker**

The software package used to provide AMQP messaging capabilities within Compute. Default package is RabbitMQ.

**multi-host**

High-availability mode for legacy (nova) networking. Each compute node handles NAT and DHCP and acts as a gateway for all of the VMs on it. A networking failure on one compute node doesn't affect VMs on other compute nodes.

**Network Address Translation (NAT)**

The process of modifying IP address information while in-transit. Supported by Compute and Networking.

**Network Time Protocol (NTP)**

A method of keeping a clock for a host or node correct through communications with a trusted, accurate time source.

**Networking**

A core OpenStack project that provides a network connectivity abstraction layer to OpenStack Compute. The project name of Networking is neutron.





trove

OpenStack project that provides database services to applications.

user

In Identity Service, each user is associated with one or more tenants, and in Compute can be associated with roles, projects, or both.

virtual machine (VM)

An operating system instance that runs on top of a hypervisor. Multiple VMs can run at the same time on the same physical host.

virtual networking

A generic term for virtualization of network functions such as switching, routing, load balancing, and security using a combination of VMs and overlays on physical network infrastructure.

Virtual Network Computing (VNC)

Open source GUI and CLI tools used for remote console access to VMs. Supported by Compute.

virtual private network (VPN)

Provided by Compute in the form of cloudpipes, specialized instances that are used to create VPNs on a per-project basis.