

# OpenStack Installation Guide

for Debian 7.0 (Wheezy)

icehouse (September 12, 2014)

Icehouse





## 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 .....	21
3. Configure OpenStack with debconf .....	22
debconf concepts .....	22
Configure the database with dbconfig-common .....	23
RabbitMQ credentials parameters .....	27
Services and the [keystone_auth token] .....	28
Register API endpoints .....	29
4. Configure the Identity Service .....	32
Identity Service concepts .....	32
Install the Identity Service .....	34
Verify the Identity Service installation .....	36
5. Install and configure the OpenStack clients .....	38
Overview .....	38
Install the OpenStack command-line clients .....	39
Set environment variables using the OpenStack RC file .....	41
Create openrc.sh files .....	42
6. Configure the Image Service .....	43
Image Service overview .....	43
Install the Image Service .....	44
Verify the Image Service installation .....	44
7. Configure Compute services .....	48
Compute service .....	48
Install Compute controller services .....	51
Configure a compute node .....	51
8. Add a networking service .....	54
OpenStack Networking (neutron) .....	54
Legacy networking (nova-network) .....	71
Next steps .....	73
9. Add the dashboard .....	74
System requirements .....	74
Install the dashboard .....	75
Next steps .....	76
10. Add the Block Storage service .....	77
Block Storage .....	77







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





# 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
<a href="#">Dashboard</a>	<a href="#">Horizon</a>	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.
<a href="#">Compute</a>	<a href="#">Nova</a>	Manages the lifecycle of compute instances in an OpenStack environment. Responsibilities include spawning, scheduling and decommissioning of virtual machines on demand.
<a href="#">Networking</a>	<a href="#">Neutron</a>	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		
<a href="#">Object Storage</a>	<a href="#">Swift</a>	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.
<a href="#">Block Storage</a>	<a href="#">Cinder</a>	Provides persistent block storage to running instances. Its pluggable driver architecture facilitates the creation and management of block storage devices.
Shared services		
<a href="#">Identity service</a>	<a href="#">Keystone</a>	Provides an authentication and authorization service for other OpenStack services. Provides a catalog of endpoints for all OpenStack services.
<a href="#">Image Service</a>	<a href="#">Glance</a>	Stores and retrieves virtual machine disk images. OpenStack Compute makes use of this during instance provisioning.
<a href="#">Telemetry</a>	<a href="#">Ceilometer</a>	Monitors and meters the OpenStack cloud for billing, benchmarking, scalability, and statistical purposes.

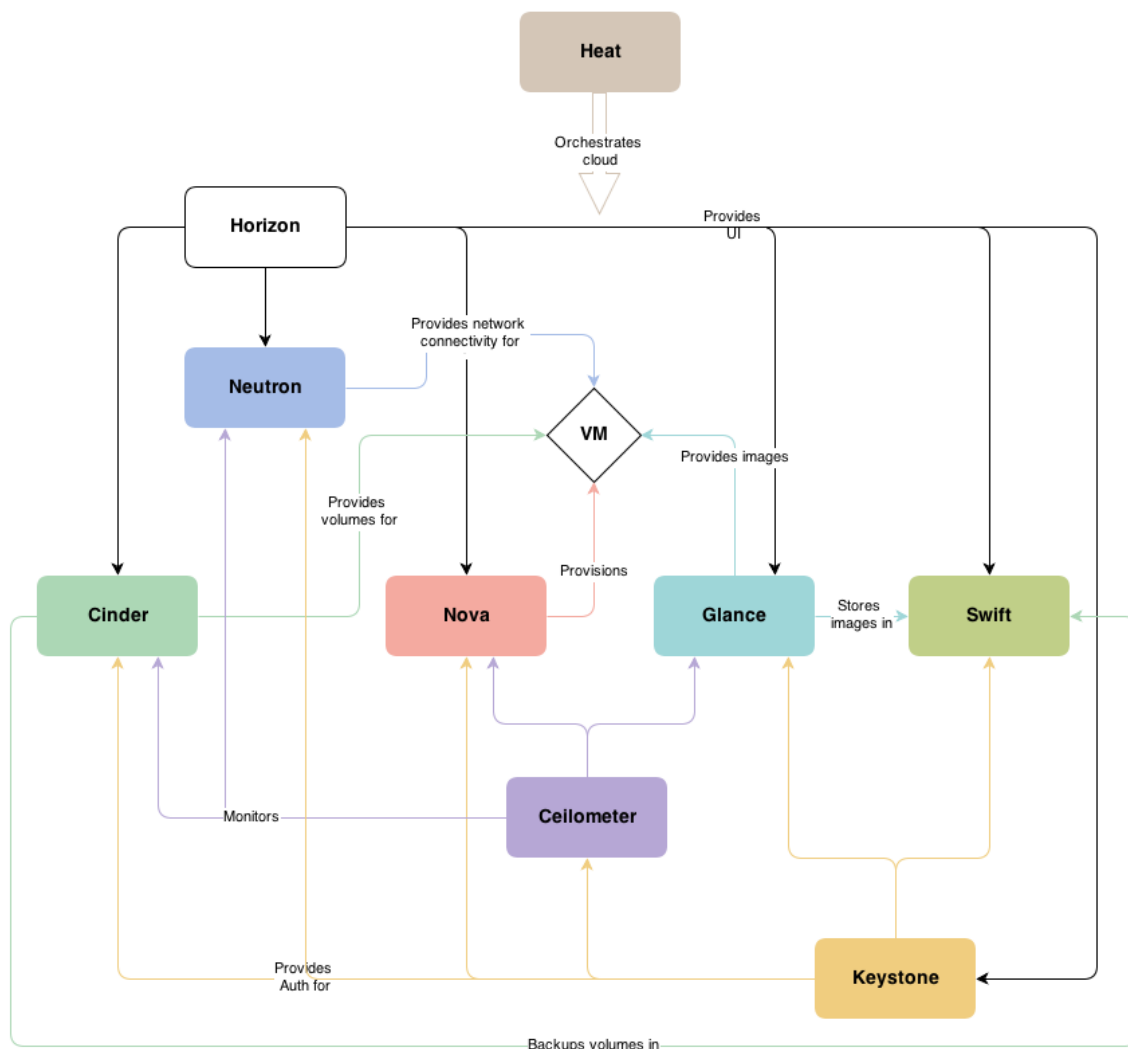
Service	Project name	Description
Higher-level services		
<i>Orchestration</i>	<i>Heat</i>	Orchestrates multiple composite cloud applications by using either the native <i>HOT</i> template format or the AWS CloudFormation template format, through both an OpenStack-native REST API and a CloudFormation-compatible 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 \(nova-network\)” \[71\]](#) in [Chapter 8, “Add a networking service” \[54\]](#). 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 \(neutron\)” \[54\]](#) in [Chapter 8, “Add a networking service” \[54\]](#). 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 .....	21

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” \[82\]](#). 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](#).

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)**



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

# network
10.0.0.21      network

# compute1
10.0.0.31      compute1
```



### Warning

You must remove or comment the line beginning with `127.0.1.1`.

## Network node

### To configure networking:

1. Configure the first interface as the management interface:  
  
IP address: 10.0.0.21  
  
Network mask: 255.255.255.0 (or /24)  
  
Default gateway: 10.0.0.1
2. Configure the second interface as the instance tunnels interface:  
  
IP address: 10.0.1.21  
  
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/network/interfaces` file to contain the following:

```
# The external network interface
auto INTERFACE_NAME
iface INTERFACE_NAME inet manual
    up ip link set dev $IFACE up
    down ip link set dev $IFACE down
```

4. Restart networking:

```
# service networking stop && service networking start
```

### 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
```



### Warning

You must remove or comment the line beginning with `127.0.1.1`.

## 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 `compute1`.
2. Edit the `/etc/hosts` file to contain the following:



```
# 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
```



### Warning

You must remove or comment the line beginning with `127.0.1.1`.

## 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/network/interfaces` file to contain the following:

```
# The external network interface
auto INTERFACE_NAME
iface INTERFACE_NAME inet manual
    up ip link set dev $IFACE up
    down ip link set dev $IFACE down
```

3. Restart networking:

```
# service networking stop && service networking start
```

### 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
```



### Warning

You must remove or comment the line beginning with `127.0.1.1`.

## 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 *compute* node:

```
# ping -c 4 computel
PING computel (10.0.0.31) 56(84) bytes of data.
64 bytes from computel (10.0.0.31): icmp_seq=1 ttl=64 time=0.263 ms
64 bytes from computel (10.0.0.31): icmp_seq=2 ttl=64 time=0.202 ms
64 bytes from computel (10.0.0.31): icmp_seq=3 ttl=64 time=0.203 ms
64 bytes from computel (10.0.0.31): icmp_seq=4 ttl=64 time=0.202 ms

--- computel 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 *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:

```
# apt-get install ntp
```

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

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.

```
# apt-get install python-mysqldb mysql-server
```



## Note

When you install the server package, you are prompted for the root password for the database. Choose a strong password and remember it.

The MySQL configuration requires some changes to work with OpenStack.

- Edit the `/etc/mysql/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
```

Restart the MySQL service to apply the changes:

```
# service mysql restart
```

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
```

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:

```
# apt-get install python-mysqldb
```

## 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.

### To use the Debian Wheezy backports archive for Icehouse

The *Icehouse* release is available only in Debian Sid (otherwise called Unstable). However, the Debian maintainers of OpenStack also maintain a non-official Debian repository for OpenStack containing Wheezy backports.

1. Install the Debian Wheezy backport repository Icehouse:

```
# echo "deb http://archive.gplhost.com/debian icehouse-backports main" >>/etc/apt/sources.list
```

2. Install the Debian Wheezy OpenStack repository for Icehouse:

```
# echo "deb http://archive.gplhost.com/debian icehouse main" >>/etc/apt/sources.list
```

3. Update the repository database and install the key:

```
# apt-get update && apt-get install gplhost-archive-keyring
```

4. Update the package database, upgrade your system, and reboot for all changes to take effect:

```
# apt-get update && apt-get dist-upgrade  
# reboot
```

Numerous archive.gplhost.com mirrors are available around the world. All are available with both FTP and HTTP protocols (you should use the closest mirror). The list of mirrors is available at <http://archive.gplhost.com/readme.mirrors>.

## Manually install python-argparse

The Debian OpenStack packages are maintained on Debian Sid (also known as Debian Unstable) - the current development version. Backported packages run correctly on Debian Wheezy with one caveat:

All OpenStack packages are written in Python. Wheezy uses Python 2.6 and 2.7, with Python 2.6 as the default interpreter; Sid has only Python 2.7. There is one packaging change between these two. In Python 2.6, you installed the `python-argparse` package separately. In Python 2.7, this package is installed by default. Unfortunately, in Python 2.7, this package does not include `Provides: python-argparse` directive.

Because the packages are maintained in Sid where the `Provides: python-argparse` directive causes an error, and the Debian OpenStack maintainer wants to maintain one version of the OpenStack packages, you must manually install the `python-argparse` on each OpenStack system that runs Debian Wheezy before you install the other OpenStack packages. Use the following command to install the package:

```
# apt-get install python-argparse
```

This caveat applies to most OpenStack packages in Wheezy.

## 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 supports several message brokers including RabbitMQ, Qpid, and ZeroMQ. However, most distributions that package OpenStack support a particular message broker. This guide covers the message broker supported by each distribution. If you prefer to implement a different message broker, consult the documentation associated with it.

- [RabbitMQ](#)
- [Qpid](#)
- [ZeroMQ](#)

### To install the message broker service

- Ubuntu and Debian use RabbitMQ.

```
# apt-get install rabbitmq-server
```

### To configure the message broker service

- The message broker creates a default account that uses `guest` for the username and password. To simplify installation of your test environment, we recommend that you use this account, but change the password for it.

Run the following command:

Replace `RABBIT_PASS` with a suitable password.

```
# rabbitmqctl change_password guest RABBIT_PASS
```

You must configure the `rabbit_password` key in the configuration file for each OpenStack service that uses the message broker.



#### Note

For production environments, you should create a unique account with suitable password. For more information on securing the message broker, see the [documentation](#).

If you decide to create a unique account with suitable password for your test environment, you must configure the `rabbit_userid` and `rabbit_password` keys in the configuration file of each OpenStack service that uses the message broker.

Congratulations, now you are ready to install OpenStack services!

## 3. Configure OpenStack with debconf

### Table of Contents

debconf concepts .....	22
Configure the database with dbconfig-common .....	23
RabbitMQ credentials parameters .....	27
Services and the [keystone_authtoken] .....	28
Register API endpoints .....	29

### debconf concepts

This chapter explains how to use the Debian `debconf` and `dbconfig-common` packages to configure OpenStack services. These packages enable users to perform configuration tasks. When users install OpenStack packages, `debconf` prompts the user for responses, which seed the contents of configuration files associated with that package. After package installation, users can update the configuration of a package by using the `dpkg-reconfigure` program.

If you are familiar with these packages and pre-seeding, you can proceed to [Chapter 4, “Configure the Identity Service” \[32\]](#).

### The Debian packages

The rules described here are from the *Debian Policy Manual*. If any rule described in this chapter is not respected, you have found a serious bug that must be fixed.

When you install or upgrade a Debian package, all configuration file values are preserved. Using the `debconf` database as a registry is considered a bug in Debian. If you edit something in any OpenStack configuration file, the `debconf` package reads that value when it prepares to prompt the user. For example, to change the log in name for the RabbitMQ messaging queue for a service, you can edit its value in the corresponding configuration file.

To opt out of using the `debconf` package, run the `dpkg-reconfigure` command and select non-interactive mode:

```
# dpkg-reconfigure -plow debconf
```

Then, `debconf` does not prompt you.

Another way to disable the `debconf` package is to prefix the `apt` command with `DEBIAN_FRONTEND=noninteractive`, as follows:

```
# DEBIAN_FRONTEND=noninteractive apt-get install nova-api
```

If you configure a package with `debconf` incorrectly, you can re-configure it, as follows:

```
# dpkg-reconfigure PACKAGE-NAME
```

This calls the post-installation script for the `PACKAGE-NAME` package after the user responds to all prompts. If you cannot install a Debian package in a non-interactive way, you have found a release-critical bug in Debian. Report it to the Debian bug tracking system.

Generally, the `-common` packages install the configuration files. For example, the `glance-common` package installs the `glance-api.conf` and `glance-registry.conf` files. So, for the Image Service, you must re-configure the `glance-common` package. The same applies for `cinder-common`, `nova-common`, and `heat-common` packages.

In `debconf`, the higher the priority for a screen, the greater the chance that the user sees that screen. If a `debconf` screen has `medium` priority and you configure the Debian system to show only `critical` prompts, which is the default in Debian, the user does not see that `debconf` screen. Instead, the default for the related package is used. In the Debian OpenStack packages, a number of `debconf` screens are set with `medium` priority. Consequently, if you want to respond to all `debconf` screens from the Debian OpenStack packages, you must run the following command and select the `medium` priority before you install any packages:

```
# dpkg-reconfigure debconf
```



### Note

The packages do not require pre-dependencies. If `dbconfig-common` is already installed on the system, the user sees all prompts. However, you cannot define the order in which the `debconf` screens appear. The user must make sense of it even if the prompts appear in an illogical order.

## Pre-seed debconf prompts

You can pre-seed all `debconf` prompts. To pre-seed means to store responses in the `debconf` database so that `debconf` does not prompt the user for responses. Pre-seeding enables a hands-free installation for users. The package maintainer creates scripts that automatically configure the services.

The following example shows how to pre-seed an automated MySQL Server installation:

```
MYSQL_PASSWORD=MYSQL_PASSWORD
echo "mysql-server-5.5 mysql-server/root_password password ${MYSQL_PASSWORD}"
mysql-server-5.5 mysql-server/root_password seen true
mysql-server-5.5 mysql-server/root_password_again password ${MYSQL_PASSWORD}
mysql-server-5.5 mysql-server/root_password_again seen true
" | debconf-set-selections
DEBIAN_FRONTEND=noninteractive apt-get install -y --force-yes mysql-server
```

The `seen true` option tells `debconf` that a specified screen was already seen by the user so do not show it again. This option is useful for upgrades.

## Configure the database with dbconfig-common

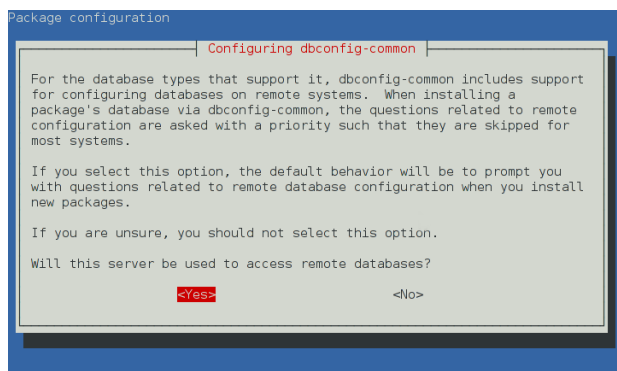
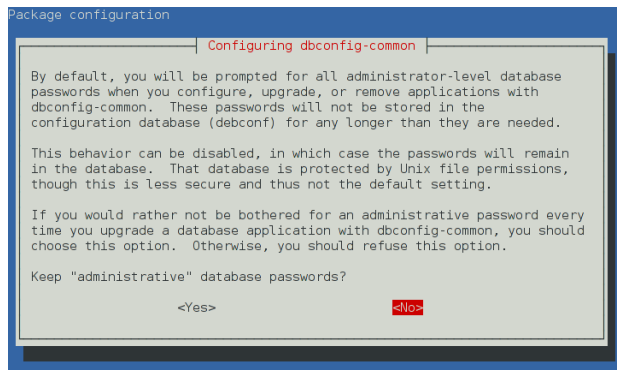
The `dbconfig-common` package provides a standard Debian interface that enables you to configure Debian database parameters. It includes localized prompts for many languages and it supports the OpenStack database back ends: SQLite, MySQL, and PostgreSQL.

By default, the `dbconfig-common` package configures the OpenStack services to use SQLite3. So if you use `debconf` in non-interactive mode and without pre-seeding, the OpenStack services that you install use SQLite3.

By default, `dbconfig-common` does not provide access to database servers over a network. If you want the `dbconfig-common` package to prompt for remote database servers that are accessed over a network and not through a UNIX socket file, reconfigure it, as follows:

```
# apt-get install dbconfig-common && dpkg-reconfigure dbconfig-common
```

These screens appear when you re-configure the `dbconfig-common` package:



Unlike other `debconf` prompts, you cannot pre-seed the responses for the `dbconfig-common` prompts by using `debconf-set-selections`. Instead, you must create a file in `/etc/dbconfig-common`. For example, you might create a keystone configuration file for `dbconfig-common` that is located in `/etc/dbconfig-common/keystone.conf`, as follows:

```
dbc_install='true'
dbc_upgrade='true'
dbc_remove=''
dbc_dbtype='mysql'
dbc_dbuser='keystone'
dbc_dbpass='MyKeyStoneDbPassWord'
dbc_dbserver=''
dbc_dbport=''
dbc_dbname='keystonedb'
dbc_dbadmin='root'
dbc_basepath=''
dbc_ssl=''
dbc_authmethod_admin=''
dbc_authmethod_user=''
```


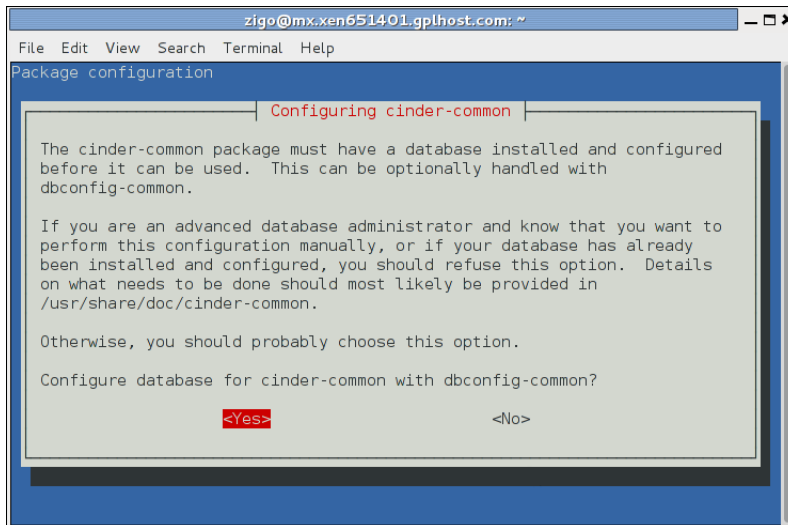





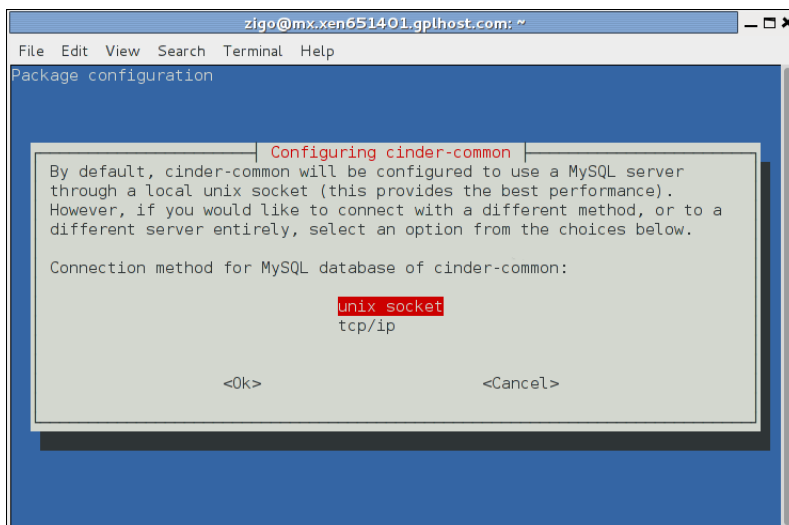
After you create this file, run this command:

```
# apt-get install keystone
```

The Identity Service is installed with MySQL as the database back end, keystone\_db as database name, and the localhost socket file.

The cinder-common package displays these screens:



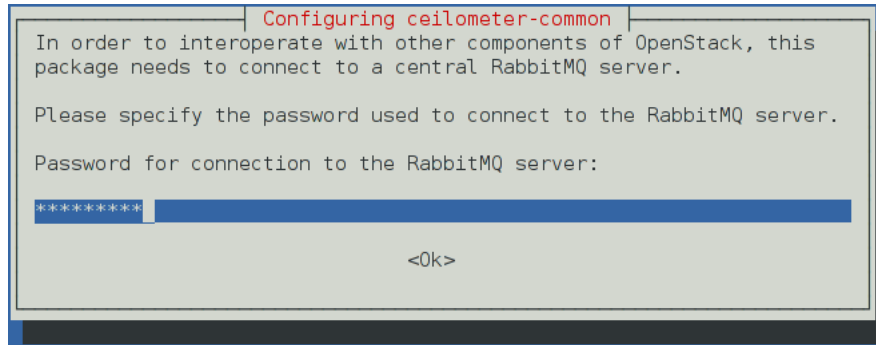


By default in Debian, you can access the MySQL server from either localhost through the socket file or 127.0.0.1. To access it over the network, you must edit the `/etc/mysql/my.cnf` file, and the `mysql.user` table. To do so, Debian provides a helper script in the `openstack-deploy` package. To use it, install the package and run:

```
# /usr/share/openstack-deploy/mysql-remote-root
```

Alternatively, if you do not want to install this package, run this script to enable remote root access:





These debconf screens appear in: `ceilometer-common`, `cinder-common`, `glance-common`, `heat-common`, `neutron-common` and `nova-common`.

## Services and the [keystone\_authtoken]



Because most OpenStack services must access the Identity Service, you must configure the IP address of the `keystone` server to be able to access it. You must also configure the `admin_tenant_name`, `admin_user`, and `admin_password` options for each service to work.

Generally, this section looks like this:

```
[keystone_authtoken]
auth_host = 127.0.0.1
auth_port = 35357
auth_protocol = http
admin_tenant_name = %SERVICE_TENANT_NAME%
admin_user = %SERVICE_USER%
admin_password = %SERVICE_PASSWORD%
```

The debconf system helps users configure the `auth_host`, `admin_tenant_name`, `admin_user` and `admin_password` options.

The following screens show an example Image Service configuration:









## 4. Configure the Identity Service

### Table of Contents

Identity Service concepts .....	32
Install the Identity Service .....	34
Verify the Identity Service installation .....	36

### 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:

<i>User</i>	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.
<i>Credentials</i>	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.
<i>Authentication</i>	<p>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.</p> <p>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.</p>
<i>Token</i>	<p>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.</p> <p>While the Identity Service supports token-based authentication in this release, the intention is for it to support additional protocols in the future. The</p>



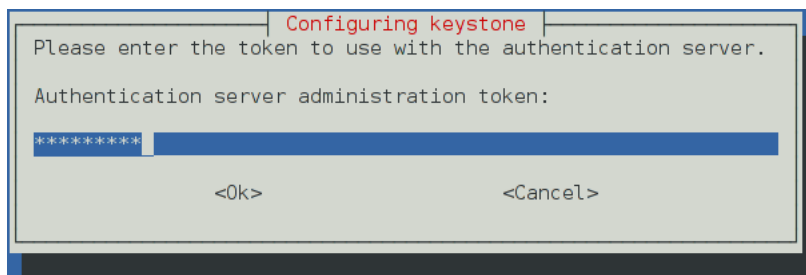


## Install the Identity Service

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

```
# apt-get install keystone
```

2. Answer to the `debconf` and `dbconfig-common` questions for setting-up the database.
3. Define an authorization token to use as a shared secret between the Identity Service and other OpenStack services. Respond to the `debconf` prompt with the value in the `admin_token` directive in the `keystone.conf` file. Use the `openssl rand -hex 10` command to generate this password.

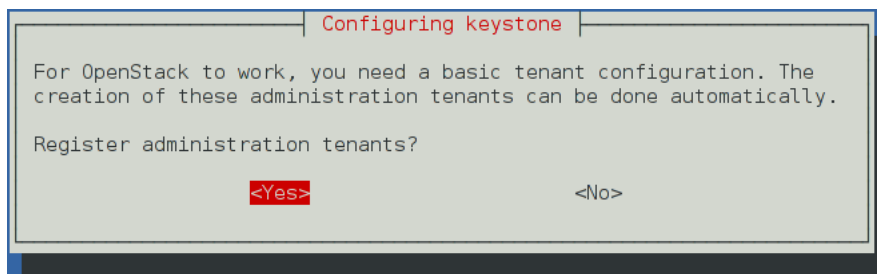


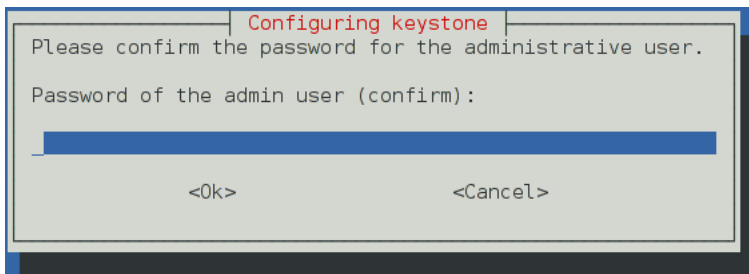
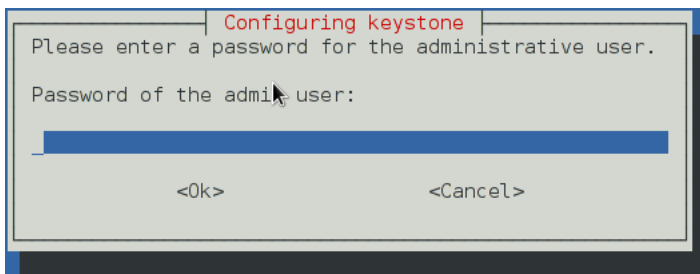
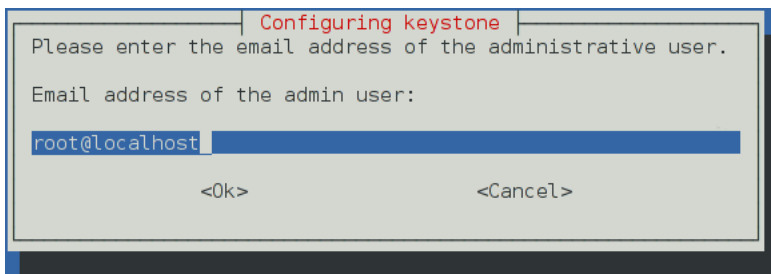
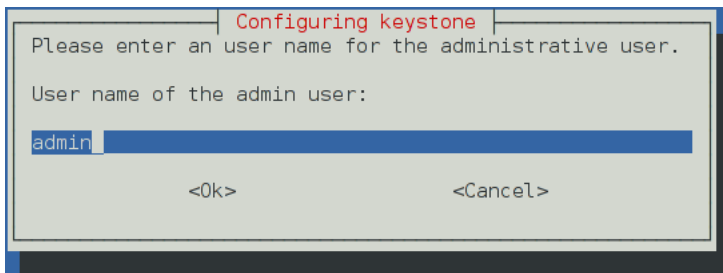
Later, you can verify that the `/etc/keystone/keystone.conf` file contains the password you have set using `debconf`:

```
[DEFAULT]
# A "shared secret" between keystone and other openstack services
admin_token = ADMIN_TOKEN
...
```

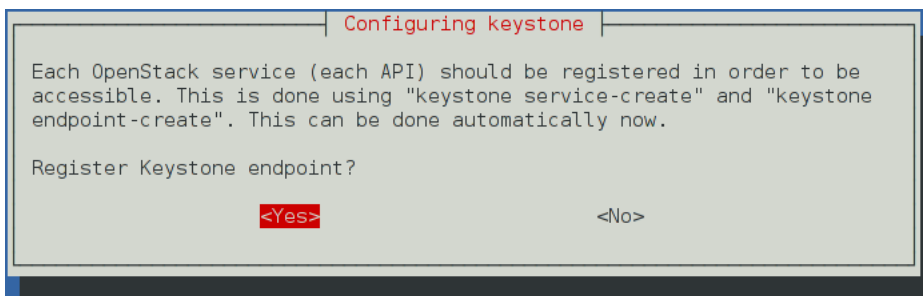
If you omit a password (for example by pressing `Enter` at the `debconf` prompt, or installing Keystone using the `Debconf` non-interactive mode) the package generates a random `ADMIN_TOKEN` value.

4. Respond to the prompts to create an administrative tenant:





5. If this is the first time you have installed the Identity Service, register the Identity Service in the service catalog:



- 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/crontabs/keystone
```

## Verify the Identity Service installation

- 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.

- 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.

- 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.

- 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
```

- Source this file to read in the environment variables:

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

- 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.

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

```
$ keystone user-list
```

id	name	enabled	email
afea5bde3be9413dbd60e479fddf9228	admin	True	admin@example.com
32acalf9a47540c29d6988091f76c934	demo	True	demo@example.com

```
$ keystone user-role-list --user admin --tenant admin
```

id	tenant_id	name	user_id
9fe2ff9ee4384b1894a90878d3e92bab		_member_	
afea5bde3be9413dbd60e479fddf9228		e519b772cb43474582fa303da62559e5	
5d3b60b66f1f438b80eaae41a77b5951		admin	
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 5, "Install and configure the OpenStack clients" \[38\]](#).

# 5. Install and configure the OpenStack clients

## Table of Contents

Overview .....	38
Install the OpenStack command-line clients .....	39
Set environment variables using the OpenStack RC file .....	41
Create openrc.sh files .....	42

The following sections contain information about working with the OpenStack clients. Recall: in the previous section, you used the **keystone** client.

You must install the client tools to complete the rest of the installation.

Configure the clients on your desktop rather than on the server so that you have a similar experience to your users.

## Overview

You can use the OpenStack command-line clients to run simple commands that make API calls. You can run these commands from the command line or in scripts to automate tasks. If you provide OpenStack credentials, you can run these commands on any computer.

Internally, each client command runs cURL commands that embed API requests. The OpenStack APIs are RESTful APIs that use the HTTP protocol, including methods, URIs, media types, and response codes.

These open-source Python clients run on Linux or Mac OS X systems and are easy to learn and use. Each OpenStack service has its own command-line client. On some client commands, you can specify a **debug** parameter to show the underlying API request for the command. This is a good way to become familiar with the OpenStack API calls.

The following table lists the command-line client for each OpenStack service with its package name and description.

**Table 5.1. OpenStack services and clients**

Service	Client	Package	Description
Block Storage	<b>cinder</b>	python-cinderclient	Create and manage volumes.
Compute	<b>nova</b>	python-novaclient	Create and manage images, instances, and flavors.
Database Service	<b>trove</b>	python-troveclient	Create and manage databases.
Identity	<b>keystone</b>	python-keystoneclient	Create and manage users, tenants, roles, endpoints, and credentials.
Image Service	<b>glance</b>	python-glanceclient	Create and manage images.

Service	Client	Package	Description
Networking	<b>neutron</b>	python-neutronclient	Configure networks for guest servers. This client was previously called <b>quantum</b> .
Object Storage	<b>swift</b>	python-swiftclient	Gather statistics, list items, update metadata, and upload, download, and delete files stored by the Object Storage service. Gain access to an Object Storage installation for ad hoc processing.
Orchestration	<b>heat</b>	python-heatclient	Launch stacks from templates, view details of running stacks including events and resources, and update and delete stacks.
Telemetry	<b>ceilometer</b>	python-ceilometerclient	Create and collect measurements across OpenStack.

An OpenStack **common** client is in development.

## Install the OpenStack command-line clients

Install the prerequisite software and the Python package for each OpenStack client.

### Install the prerequisite software

The following table lists the software that you need to have to run the command-line clients, and provides installation instructions as needed.

**Table 5.2. Prerequisite software**

Prerequisite	Description
Python 2.6 or later	Currently, the clients do not support Python 3.
setuptools package	<p>Installed by default on Mac OS X.</p> <p>Many Linux distributions provide packages to make setuptools easy to install. Search your package manager for setuptools to find an installation package. If you cannot find one, download the setuptools package directly from <a href="http://pypi.python.org/pypi/setuptools">http://pypi.python.org/pypi/setuptools</a>.</p> <p>The recommended way to install setuptools on Microsoft Windows is to follow the documentation provided <a href="#">on the setuptools website</a>. Another option is to use the unofficial binary installer maintained by Christoph Gohlke (<a href="http://www.lfd.uci.edu/~gohlke/pythonlibs/#setuptools">http://www.lfd.uci.edu/~gohlke/pythonlibs/#setuptools</a>).</p>
pip package	<p>To install the clients on a Linux, Mac OS X, or Microsoft Windows system, use pip. It is easy to use, ensures that you get the latest version of the clients from the <a href="#">Python Package Index</a>, and lets you update or remove the packages later on.</p> <p>Install pip through the package manager for your system:</p> <p><b>MacOS.</b></p> <pre># easy_install pip</pre> <p><b>Microsoft Windows.</b> Ensure that the <code>C:\Python27\Scripts</code> directory is defined in the <code>PATH</code> environment variable, and use the <code>easy_install</code> command from the setuptools package:</p> <pre>C:\&gt;easy_install pip</pre> <p>Another option is to use the unofficial binary installer provided by Christoph Gohlke (<a href="http://www.lfd.uci.edu/~gohlke/pythonlibs/#pip">http://www.lfd.uci.edu/~gohlke/pythonlibs/#pip</a>).</p> <p><b>Ubuntu and Debian.</b></p>

Prerequisite	Description
	<pre># apt-get install python-pip</pre>
	Red Hat Enterprise Linux, CentOS, or Fedora. A packaged version available in <a href="#">RDO</a> enables you to use yum to install the clients, or you can install pip and use it to manage client installation:
	<pre># yum install python-pip</pre>
	openSUSE 12.2 and earlier. A <a href="#">packaged version available in the Open Build Service</a> enables you to use rpm or zypper to install the clients, or you can install pip and use it to manage client installation:
	<pre># zypper install python-pip</pre>
	openSUSE 12.3 and later. A packaged version enables you to use rpm or zypper to install the clients. See <a href="#">the section called "Install the clients" [40]</a>

## 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” \[41\]](#), 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
```

## 6. Configure the Image Service

### Table of Contents

Image Service overview .....	43
Install the Image Service .....	44
Verify the Image Service installation .....	44

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.



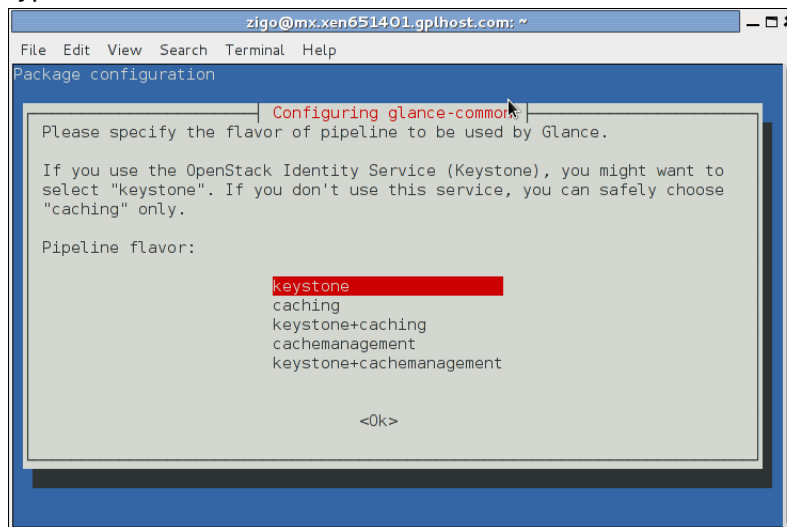
### Note

This procedure assumes you set the appropriate environment variables to your credentials as described in [the section called "Verify the Identity Service installation" \[36\]](#).

1. Install the Image Service on the controller node:

```
# apt-get install glance python-glanceclient
```

2. Respond to prompts for [database management](#), [\[keystone\\_auth token\] settings](#), [RabbitMQ credentials](#) and [API endpoint](#) registration. You must also select the caching type:



## 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.



```

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

```

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

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

Alternatively, the upload to the Image Service can be done without having to use local disk space to store the file, by use of the `--copy-from` parameter.

For example:

```
$ glance image-create --name="cirros-0.3.2-x86_64" --disk-format=qcow2 \
--container-format=bare --is-public=true \
--copy-from http://cdn.download.cirros-cloud.net/0.3.2/cirros-0.3.2-x86_64-disk.img
```

```

+-----+-----+
| Property | Value |
+-----+-----+
| checksum | 64d7c1cd2b6f60c92c14662941cb7913 |
| container_format | bare |
| created_at | 2014-04-08T06:13:18 |
| deleted | False |
| disk_format | qcow2 |
+-----+-----+

```



# 7. Configure Compute services

## Table of Contents

Compute service .....	48
Install Compute controller services .....	51
Configure a compute node .....	51

## 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.



- `nova-conductor` module. Mediates interactions between `nova-compute` and the database. Aims to eliminate direct accesses to the cloud database made by `nova-compute`. The `nova-conductor` module scales horizontally. However, do not deploy it on any nodes where `nova-compute` runs. For more information, see [A new Nova service: nova-conductor](#).

### Networking for VMs

- `nova-network` worker daemon. Similar to `nova-compute`, it accepts networking tasks from the queue and performs tasks to manipulate the network, such as setting up bridging interfaces or changing iptables rules. This functionality is being migrated to OpenStack Networking, which is a separate OpenStack service.
- `nova-dhcpbridge` script. Tracks IP address leases and records them in the database by using the `dnsmasq dhcp-script` facility. This functionality is being migrated to OpenStack Networking. OpenStack Networking provides a different script.

## Console interface

- `nova-consoleauth` daemon. Authorizes tokens for users that console proxies provide. See `nova-novncproxy` and `nova-xvpncproxy`. 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-xvpncproxy` 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.

In Debian, a unique `nova-consoleproxy` package provides the `nova-novncproxy`, `nova-spicehtml5proxy`, and `nova-xvpncproxy` packages. To select packages, edit the `/etc/default/nova-consoleproxy` file or use the debconf interface. You can also manually edit the `/etc/default/nova-consoleproxy` file and stop and start the console daemons.

## 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.

```
# apt-get install nova-api nova-cert nova-conductor nova-consoleauth \  
nova-novncproxy nova-scheduler python-novaclient
```

2. Respond to the prompts for [database management](#), [\[keystone\\_auth token\] settings](#), [RabbitMQ credentials](#), and [API endpoint](#) registration. The **nova-manage db sync** command runs automatically.
3. Set the `my_ip`, `vncserver_listen`, and `vncserver_proxyclient_address` configuration options to the management interface IP address of the controller node:

In Debian, the `debconf` package automatically sets up `my_ip` parameter but you must edit the `/etc/nova/nova.conf` file to configure the `vncserver_listen` and `vncserver_proxyclient_address` options, which appear at the end of the file:

```
[DEFAULT]  
...  
my_ip = 10.0.0.11  
vncserver_listen = 10.0.0.11  
vncserver_proxyclient_address = 10.0.0.11
```

4. To verify your configuration, list available images:

```
$ nova image-list  
+-----+-----+-----+  
+-----+ | Name | Status |  
| ID | | |  
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:

```
# apt-get install nova-compute-kvm
```



### Note

To use the meta-packages and install other components on your compute node, such as OpenStack Networking and Ceilometer agents, run this command:

```
# apt-get install openstack-compute-node
```

The controller node has the `openstack-proxy-node` and `openstack-toaster` meta-packages that install `openstack-proxy-node` and `openstack-toaster` at the same time.

2. Respond to the prompts for [database management](#), [\[keystone\\_authtoken\] settings](#), [RabbitMQ credentials](#), and [API endpoint](#) registration.
3. Configure Compute to provide remote console access to instances.

Edit `/etc/nova/nova.conf` and add the following keys under the `[DEFAULT]` section:

```
[DEFAULT]
...
my_ip = 10.0.0.31
vnc_enabled = True
vncserver_listen = 0.0.0.0
vncserver_proxyclient_address = 10.0.0.31
novncproxy_base_url = http://controller:6080/vnc_auto.html
```

4. Specify the host that runs the Image Service. Edit `/etc/nova/nova.conf` file and add these lines to the `[DEFAULT]` section:

```
[DEFAULT]
...
glance_host = controller
```

5. 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.

- Edit the `[libvirt]` section in the `/etc/nova/nova-compute.conf` file to modify this key:

```
[libvirt]
...
virt_type = qemu
```



## 8. Add a networking service

### Table of Contents

OpenStack Networking (neutron) .....	54
Legacy networking (nova-network) .....	71
Next steps .....	73

Configuring networking in OpenStack can be a bewildering experience. This guide provides step-by-step instructions for both OpenStack Networking (neutron) and the legacy networking (nova-network) service. If you are unsure which to use, we recommend trying OpenStack Networking because it offers a considerable number of features and flexibility including *plug-ins* for a variety of emerging products supporting *virtual networking*. See the [Networking](#) chapter of the *OpenStack Cloud Administrator Guide* for more information.

## OpenStack Networking (neutron)

### Networking concepts

OpenStack Networking (neutron) manages all of the networking facets for the Virtual Networking Infrastructure (VNI) and the access layer aspects of the Physical Networking Infrastructure (PNI) in your OpenStack environment. OpenStack Networking allows tenants to create advanced virtual network topologies including services such as *firewalls*, *load balancers*, and *virtual private networks (VPNs)*.

Networking provides the following object abstractions: networks, subnets, and routers. Each has functionality that mimics its physical counterpart: networks contain subnets, and routers route traffic between different subnet and networks.

Any given Networking set up has at least one external network. This network, unlike the other networks, is not merely a virtually defined network. Instead, it represents the view into a slice of the external network that is accessible outside the OpenStack installation. IP addresses on the Networking external network are accessible by anybody physically on the outside network. Because this network merely represents a slice of the outside network, DHCP is disabled on this network.

In addition to external networks, any Networking set up has one or more internal networks. These software-defined networks connect directly to the VMs. Only the VMs on any given internal network, or those on subnets connected through interfaces to a similar router, can access VMs connected to that network directly.

For the outside network to access VMs, and vice versa, routers between the networks are needed. Each router has one gateway that is connected to a network and many interfaces that are connected to subnets. Like a physical router, subnets can access machines on other subnets that are connected to the same router, and machines can access the outside network through the gateway for the router.

Additionally, you can allocate IP addresses on external networks to ports on the internal network. Whenever something is connected to a subnet, that connection is called a port. 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

#### To install the Networking components

- ```
# apt-get install neutron-server
```



#### Note

Debian does not use a separate ML2 plug-in package.

#### 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. Respond to prompts for [database management](#), [Identity service credentials](#), [service endpoint registration](#), and [message broker credentials](#).
2. During the installation, you will also be prompted for which Networking plug-in to use. This will automatically fill the `core_plugin` directive in the `/etc/neutron/neutron.conf` file.





| Property    | Value                            |
|-------------|----------------------------------|
| description | Service Tenant                   |
| enabled     | True                             |
| id          | f727b5ec2ceb4d71bad86dfc414449bf |
| name        | service                          |

## 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. However, the controller node does not need the OVS agent or service because it does not handle instance network traffic.

- Edit the `/etc/neutron/plugins/ml2/ml2_conf.ini` file:

Add the following keys to the `[ml2]` section:

```
[ml2]
...
type_drivers = gre
tenant_network_types = gre
mechanism_drivers = openvswitch
```

Add the following key to the `[ml2_type_gre]` section:

```
[ml2_type_gre]
...
tunnel_id_ranges = 1:1000
```

Add the `[securitygroup]` section and the following keys to it:

```
[securitygroup]
...
firewall_driver = neutron.agent.linux.iptables_firewall.
OVSHybridIptablesFirewallDriver
enable_security_group = True
```

## 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.

- Edit the `/etc/nova/nova.conf` and add the following keys to the `[DEFAULT]` section:

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

```
[DEFAULT]
...
network_api_class = nova.network.neutronv2.api.API
neutron_url = http://controller:9696
neutron_auth_strategy = keystone
neutron_admin_tenant_name = service
neutron_admin_username = neutron
neutron_admin_password = NEUTRON_PASS
neutron_admin_auth_url = http://controller:35357/v2.0
linuxnet_interface_driver = nova.network.linux_net.LinuxOVSIInterfaceDriver
firewall_driver = nova.virt.firewall.NoopFirewallDriver
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. Restart the Compute services:

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

2. Restart the Networking service:

```
# service neutron-server restart
```

## 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

- ```
# apt-get install neutron-plugin-ml2 neutron-plugin-openvswitch-agent
  openvswitch-datapath-dkms \
  neutron-l3-agent neutron-dhcp-agent
```

## To configure the Networking common components

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

1. Respond to prompts for [database management](#), [Identity service credentials](#), [service endpoint registration](#), and [message broker credentials](#).
2. Configure Networking to use the Modular Layer 2 (ML2) plug-in and associated services:
  - Edit the `/etc/neutron/neutron.conf` file and add the following keys to the `[DEFAULT]` section:

```
[DEFAULT]
...
core_plugin = ml2
service_plugins = router
allow_overlapping_ips = True
```



### 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.

- Edit the `/etc/neutron/l3_agent.ini` file and add the following keys to the `[DEFAULT]` section:

```
[DEFAULT]
...
interface_driver = neutron.agent.linux.interface.OVSInterfaceDriver
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. Edit the `/etc/neutron/dhcp_agent.ini` file and add the following keys to the `[DEFAULT]` section:

```
[DEFAULT]
...
interface_driver = neutron.agent.linux.interface.OVSInterfaceDriver
dhcp_driver = neutron.agent.linux.dhcp.Dnsmasq
use_namespaces = True
```







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

## To finalize the installation

- Restart the Networking services:

```
# service neutron-plugin-openvswitch-agent restart
# service neutron-l3-agent restart
# service neutron-dhcp-agent restart
# service neutron-metadata-agent restart
```

## Configure compute node

### Prerequisites

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

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

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

- Implement the changes:

```
# sysctl -p
```

### To install the Networking components

- ```
# apt-get install neutron-common neutron-plugin-ml2 neutron-plugin-openvswitch-agent \
  openvswitch-datapath-dkms
```

### To configure the Networking common components

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

- Respond to prompts for [database management](#), [Identity service credentials](#), [service endpoint registration](#), and [message broker credentials](#).
- Configure Networking to use the Modular Layer 2 (ML2) plug-in and associated services:
  - Edit the `/etc/neutron/neutron.conf` file and add the following keys to the `[DEFAULT]` section:

```
[DEFAULT]
...
core_plugin = ml2
service_plugins = router
allow_overlapping_ips = True
```



### 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.

- Edit the `/etc/neutron/plugins/ml2/ml2_conf.ini` file:

Add the following keys to the `[ml2]` section:

```
[ml2]
...
type_drivers = gre
tenant_network_types = gre
mechanism_drivers = openvswitch
```

Add the following keys to the `[ml2_type_gre]` section:

```
[ml2_type_gre]
...
tunnel_id_ranges = 1:1000
```

Add the `[ovs]` section and the following keys to it:

Replace `INSTANCE_TUNNELS_INTERFACE_IP_ADDRESS` with the IP address of the instance tunnels network interface on your compute node.

```
[ovs]
...
local_ip = INSTANCE_TUNNELS_INTERFACE_IP_ADDRESS
tunnel_type = gre
enable_tunneling = True
```

Add the `[securitygroup]` section and the following keys to it:

```
[securitygroup]
...
firewall_driver = neutron.agent.linux.iptables_firewall.OVSHybridIptablesFirewallDriver
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. Restart the OVS service:

```
# service openvswitch-switch restart
```

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.







## 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:



## To create the tenant network

1. Source the demo tenant credentials:

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

2. Create the network:

```
$ neutron net-create demo-net
Created a new network:
+-----+-----+
| Field          | Value                                |
+-----+-----+
admin_state_up	True
id	ac108952-6096-4243-adf4-bb6615b3de28
name	demo-net
shared	False
status	ACTIVE
subnets	
tenant_id	cdef0071a0194d19ac6bb63802dc9bae
+-----+-----+
```

Like the external network, your tenant network also requires a subnet attached to it. You can specify any valid subnet because the architecture isolates tenant networks. Replace *TENANT\_NETWORK\_CIDR* with the subnet you want to associate with the tenant network. Replace *TENANT\_NETWORK\_GATEWAY* with the gateway you want to associate with this network, typically the ".1" IP address. By default, this subnet will use DHCP so your instances can obtain IP addresses.

## To create a subnet on the tenant network

- Create the subnet:

```
$ neutron subnet-create demo-net --name demo-subnet \
--gateway TENANT_NETWORK_GATEWAY TENANT_NETWORK_CIDR
```

Example using 192.168.1.0/24:

```
$ neutron subnet-create demo-net --name demo-subnet \
--gateway 192.168.1.1 192.168.1.0/24
Created a new subnet:
+-----+-----+
| Field          | Value                                |
+-----+-----+
allocation_pools	{"start": "192.168.1.2", "end": "192.168.1.254"}
cidr	192.168.1.0/24
dns_nameservers	
enable_dhcp	True
gateway_ip	192.168.1.1
host_routes	
+-----+-----+
```

```
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.



## Note

If you are building your OpenStack nodes as virtual machines, you must configure the hypervisor to permit promiscuous mode on the external network.

### To verify network connectivity

- Ping the tenant router gateway:

```
$ ping -c 4 203.0.113.101
PING 203.0.113.101 (203.0.113.101) 56(84) bytes of data.
64 bytes from 203.0.113.101: icmp_req=1 ttl=64 time=0.619 ms
64 bytes from 203.0.113.101: icmp_req=2 ttl=64 time=0.189 ms
64 bytes from 203.0.113.101: icmp_req=3 ttl=64 time=0.165 ms
64 bytes from 203.0.113.101: icmp_req=4 ttl=64 time=0.216 ms

--- 203.0.113.101 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 2999ms
rtt min/avg/max/mdev = 0.165/0.297/0.619/0.187 ms
```

## Legacy networking (nova-network)

### Configure controller node

Legacy networking primarily involves compute nodes. However, you must configure the controller node to use it.

#### To configure legacy networking

1. Edit the `/etc/nova/nova.conf` file and add the following keys to the `[DEFAULT]` section:

```
[DEFAULT]
...
network_api_class = nova.network.api.API
security_group_api = nova
```

2. Restart the Compute services:

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

### Configure compute node

This section covers deployment of a simple *flat network* that provides IP addresses to your instances via *DHCP*. If your environment includes multiple compute nodes, the *multi-host* feature provides redundancy by spreading network functions across compute nodes.

#### To install legacy networking components

- ```
# apt-get install nova-network nova-api-metadata
```

## To configure legacy networking

1. Edit the `/etc/nova/nova.conf` file and add the following keys to the `[DEFAULT]` section:

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

```
[DEFAULT]
...
network_api_class = nova.network.api.API
security_group_api = nova
firewall_driver = nova.virt.libvirt.firewall.IptablesFirewallDriver
network_manager = nova.network.manager.FlatDHCPManager
network_size = 254
allow_same_net_traffic = False
multi_host = True
send_arp_for_ha = True
share_dhcp_address = True
force_dhcp_release = True
flat_network_bridge = br100
flat_interface = INTERFACE_NAME
public_interface = INTERFACE_NAME
```

2. Restart the services:

```
# service nova-network restart
# service nova-api-metadata restart
```

## 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 \
```





## 9. Add the dashboard

### Table of Contents

System requirements .....	74
Install the dashboard .....	75
Next steps .....	76

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" \[74\]](#).



### 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:

```
# apt-get install apache2 memcached libapache2-mod-wsgi openstack-  
dashboard
```



### Note for Debian users

To install the Apache package:

```
# apt-get install openstack-dashboard-apache
```

This command installs and configures Apache correctly, provided that the user asks for it during the debconf prompts. The default SSL certificate is self-signed, and it is probably wise to have it signed by a root Certificate Authority (CA).

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

Open `/etc/openstack-dashboard/local_settings.py` 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/memcached.conf`.

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.py` file.

Change the following parameter: `TIME_ZONE = "UTC"`

3. Update the `ALLOWED_HOSTS` in `local_settings.py` to include the addresses you wish to access the dashboard from.

Edit `/etc/openstack-dashboard/local_settings.py`:

```
ALLOWED_HOSTS = ['localhost', 'my-desktop']
```

4. This guide assumes that you are running the Dashboard on the controller node. You can easily run the dashboard on a separate server, by changing the appropriate settings in `local_settings.py`.

Edit `/etc/openstack-dashboard/local_settings.py` and change `OPENSTACK_HOST` to the hostname of your Identity Service:

```
OPENSTACK_HOST = "controller"
```

5. Start the Apache web server and memcached:

```
# service apache2 restart  
# service memcached restart
```

6. You can now access the dashboard at `https://controller/`.

Login with credentials for any user that you created with the OpenStack Identity Service.

## Next steps

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

# 10. Add the Block Storage service

## Table of Contents

Block Storage .....	77
Configure a Block Storage service controller .....	77
Configure a Block Storage service node .....	78
Verify the Block Storage installation .....	79
Next steps .....	80

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” \[78\]](#).

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

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

```
# apt-get install cinder-api cinder-scheduler
```

2. Respond to the prompts for [database management](#), [\[keystone\\_authtoken\] settings](#), [RabbitMQ credentials](#), and [API endpoint](#) registration.

3. 

```
...
[keystone_authtoken]
auth_uri = http://controller:5000
auth_host = controller
auth_port = 35357
auth_protocol = http
admin_tenant_name = service
admin_user = cinder
admin_password = CINDER_PASS
```

## 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:

```
# apt-get install cinder-volume
```

5. Respond to the `debconf` prompts about the [database management](#), [\[keystone\\_auth token\] settings](#), and [RabbitMQ credentials](#). Make sure to enter the same details as you did for your Block Storage service controller node.

Another screen prompts you for the `volume-group` to use. The Debian package configuration script detects every active volume group and tries to use the first one it sees, provided that the `lvm2` package was installed before Block Storage. This should be the case if you configured the volume group first, as this guide recommends.

If you have only one active volume group on your Block Storage service node, you do not need to manually enter its name when you install the `cinder-volume` package because it is detected automatically. If no `volume-group` is available when you install `cinder-common`, you must use `dpkg-reconfigure` to manually configure or reconfigure `cinder-common`.

6. 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:

```
[DEFAULT]
...
glance_host = controller
```

## 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_volid	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.



# 11. Add Object Storage

## Table of Contents

Object Storage service .....	81
System requirements for Object Storage .....	82
Plan networking for Object Storage .....	82
Example of Object Storage installation architecture .....	84
Install Object Storage .....	85
Install and configure storage nodes .....	87
Install and configure the proxy node .....	88
Start services on the storage nodes .....	91
Verify the installation .....	91
Add another proxy server .....	92
Next steps .....	92

The OpenStack Object Storage services work together to provide object storage and retrieval through a REST API. For this example architecture, you must have already installed the Identity Service, also known as Keystone.

## Object Storage service

The Object Storage service is a highly scalable and durable multi-tenant object storage system for large amounts of unstructured data at low cost through a RESTful HTTP API.

It includes the following components:

- Proxy servers (`swift-proxy-server`). Accepts Object Storage API and raw HTTP requests to upload files, modify metadata, and create containers. It also serves file or container listings to web browsers. To improve performance, the proxy server can use an optional cache usually deployed with memcache.
- Account servers (`swift-account-server`). Manage accounts defined with the Object Storage service.
- Container servers (`swift-container-server`). Manage a mapping of containers, or folders, within the Object Storage service.
- Object servers (`swift-object-server`). Manage actual objects, such as files, on the storage nodes.
- A number of periodic processes. Performs housekeeping tasks on the large data store. The replication services ensure consistency and availability through the cluster. Other periodic processes include auditors, updaters, and reapers.
- Configurable WSGI middleware that handles authentication. Usually the Identity Service.

## 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 11.1. Hardware recommendations**

Server	Recommended Hardware	Notes
Object Storage object servers	Processor: dual quad core Memory: 8 or 12 GB RAM Disk space: optimized for cost per GB 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.  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 Memory: 8 or 12 GB RAM 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 Network: one 1 GB Network Interface Card (NIC)	Higher network throughput offers better performance for supporting many API requests.  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

APIs and storage network as necessary, this section offers recommendations and required minimum sizes. Throughput of at least 1000 Mbps is suggested.

This guide describes the following networks:

- A mandatory public network. Connects to the proxy server.
- A mandatory storage network. Not accessible from outside the cluster. All nodes connect to this network.
- An optional replication network. Not accessible from outside the cluster. Dedicated to replication traffic among storage nodes. Must be configured in the Ring.

This figure shows the basic architecture for the public network, the storage network, and the optional replication network.




By default, all of the OpenStack Object Storage services, as well as the rsync daemon on the storage nodes, are configured to listen on their `STORAGE_LOCAL_NET` IP addresses.

If you configure a replication network in the Ring, the Account, Container and Object servers listen on both the `STORAGE_LOCAL_NET` and `STORAGE_REPLICATION_NET` IP addresses. The rsync daemon only listens on the `STORAGE_REPLICATION_NET` IP address.

Public Network (Publicly routable IP range)

Provides public IP accessibility to the API endpoints within the cloud infrastructure.





## 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.





```
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/default/rsync`:

```
RSYNC_ENABLE=true
```

- Start the `rsync` service:

```
# service rsync 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:

```
# apt-get install swift swift-proxy memcached python-keystoneclient
python-swiftclient python-webob
```

- Modify `memcached` to listen on the default interface on a local, non-public network. Edit this line in the `/etc/memcached.conf` file:



```
-l 127.0.0.1
```

Change it to:

```
-l PROXY_LOCAL_NET_IP
```

- Restart the memcached service:

```
# service memcached restart
```

- Create `/etc/swift/proxy-server.conf`:

```
[DEFAULT]
bind_port = 8080
user = swift

[pipeline:main]
pipeline = healthcheck cache authtoken keystoneauth proxy-server

[app:proxy-server]
use = egg:swift#proxy
allow_account_management = true
account_autocreate = true

[filter:keystoneauth]
use = egg:swift#keystoneauth
operator_roles = Member,admin,swiftoperator

[filter:authtoken]
paste.filter_factory = keystoneclient.middleware.auth_token:filter_factory

# Delaying the auth decision is required to support token-less
# usage for anonymous referrers ('.r:*').
delay_auth_decision = true

# auth_* settings refer to the Keystone server
auth_protocol = http
auth_host = controller
auth_port = 35357

# the service tenant and swift username and password created in Keystone
admin_tenant_name = service
admin_user = swift
admin_password = SWIFT_PASS

[filter:cache]
use = egg:swift#memcache

[filter:catch_errors]
use = egg:swift#catch_errors

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



### Note

If you run multiple memcache servers, put the multiple IP:port listings in the `[filter:cache]` section of the `/etc/swift/proxy-server.conf` file:



```
# swift-ring-builder account.builder rebalance
# swift-ring-builder container.builder rebalance
# swift-ring-builder object.builder rebalance
```



### Note

Rebalancing rings can take some time.

9. Copy the `account.ring.gz`, `container.ring.gz`, and `object.ring.gz` files to each of the Proxy and Storage nodes in `/etc/swift`.
10. Make sure the `swift` user owns all configuration files:

```
# chown -R swift:swift /etc/swift
```

11. Restart the Proxy service:

```
# service swift-proxy restart
```

## Start services on the storage nodes

Now that the ring files are on each storage node, you can start the services. On each storage node, run the following command:

```
# for service in \
  swift-object swift-object-replicator swift-object-updater swift-object-
auditor \
  swift-container swift-container-replicator swift-container-updater swift-
container-auditor \
  swift-account swift-account-replicator swift-account-reaper swift-account-
auditor; do \
  service $service start; done
```



### Note

To start all swift services at once, run the command:

```
# swift-init all start
```

To know more about `swift-init` command, run:

```
$ man swift-init
```

## Verify the installation

You can run these commands from the proxy server or any server that has access to the Identity Service.

1. Make sure that your credentials are set up correctly in the `admin-openrc.sh` file and source it:

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

2. Run the following `swift` command:

```
$ swift stat
Account: AUTH_11b9758b7049476d9b48f7a91ea11493
Containers: 0
  Objects: 0
  Bytes: 0
Content-Type: text/plain; charset=utf-8
X-Timestamp: 1381434243.83760
X-Trans-Id: txdcdd594565214fb4a2d33-0052570383
X-Put-Timestamp: 1381434243.83760
```

3. Run the following **swift** commands to upload files to a container. Create the `test.txt` and `test2.txt` test files locally if needed.

```
$ swift upload myfiles test.txt
$ swift upload myfiles test2.txt
```

4. Run the following **swift** command to download all files from the `myfiles` container:

```
$ swift download myfiles
test2.txt [headers 0.267s, total 0.267s, 0.000s MB/s]
test.txt [headers 0.271s, total 0.271s, 0.000s MB/s]
```

## Add another proxy server

To provide additional reliability and bandwidth to your cluster, you can add proxy servers. You can set up an additional proxy node the same way that you set up the first proxy node but with additional configuration steps.

After you have more than two proxies, you must load balance them; your storage endpoint (what clients use to connect to your storage) also changes. You can select from different strategies for load balancing. For example, you could use round-robin DNS, or a software or hardware load balancer (like pound) in front of the two proxies. You can then point your storage URL to the load balancer, configure an initial proxy node and complete these steps to add proxy servers.

1. Update the list of memcache servers in the `/etc/swift/proxy-server.conf` file for added proxy servers. If you run multiple memcache servers, use this pattern for the multiple IP:port listings in each proxy server configuration file:

```
10.1.2.3:11211,10.1.2.4:11211
```

```
[filter:cache]
use = egg:swift#memcache
memcache_servers = PROXY_LOCAL_NET_IP:11211
```

2. Copy ring information to all nodes, including new proxy nodes. Also, ensure that the ring information gets to all storage nodes.
3. After you sync all nodes, make sure that the admin has keys in `/etc/swift` and the ownership for the ring file is correct.

## Next steps

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

## 12. Add the Orchestration service

### Table of Contents

Orchestration service overview .....	93
Install the Orchestration service .....	93
Verify the Orchestration service installation .....	95
Next steps .....	96

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:

```
# apt-get install heat-api heat-api-cfn heat-engine
```



## Verify the Orchestration service installation

To verify that the Orchestration service is installed and configured correctly, make sure that your credentials are set up correctly in the `demo-openrc.sh` file. Source the file, as follows:

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

The Orchestration Module uses templates to describe stacks. To learn about the template languages, see [the Template Guide](#) in the [Heat developer documentation](#).

Create a test template in the `test-stack.yml` file with the following content:

```
heat_template_version: 2013-05-23

description: Test Template

parameters:
  ImageID:
    type: string
    description: Image use to boot a server
  NetID:
    type: string
    description: Network ID for the server

resources:
  server1:
    type: OS::Nova::Server
    properties:
      name: "Test server"
      image: { get_param: ImageID }
      flavor: "m1.tiny"
      networks:
        - network: { get_param: NetID }

outputs:
  server1_private_ip:
    description: IP address of the server in the private network
    value: { get_attr: [ server1, first_address ] }
```

Use the `heat stack-create` command to create a stack from this template:

```
$ NET_ID=$(nova net-list | awk '/ demo-net / { print $2 }')
$ heat stack-create -f test-stack.yml \
  -P "ImageID=cirros-0.3.2-x86_64;NetID=$NET_ID" testStack
+-----+-----+-----+-----+
| id                | stack_name | stack_status |
| creation_time     |            |              |
+-----+-----+-----+-----+
| 477d96b4-d547-4069-938d-32ee990834af | testStack  | CREATE_IN_PROGRESS |
| 2014-04-06T15:11:01Z |            |                  |
+-----+-----+-----+-----+
```

Verify that the stack was created successfully with the `heat stack-list` command:

```
$ heat stack-list
```

```
+-----+-----+-----+
+-----+
| id                | stack_name | stack_status |
| creation_time     |            |              |
+-----+-----+-----+
| 477d96b4-d547-4069-938d-32ee990834af | testStack  | CREATE_COMPLETE |
| 2014-04-06T15:11:01Z |            |              |
+-----+-----+-----+
+-----+
```

## Next steps

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



# 13. Add the Telemetry module

## Table of Contents

Telemetry .....	97
Install the Telemetry module .....	98
Install the Compute agent for Telemetry .....	100
Configure the Image Service for Telemetry .....	101
Add the Block Storage service agent for Telemetry .....	101
Configure the Object Storage service for Telemetry .....	102
Verify the Telemetry installation .....	102
Next steps .....	104

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.



4. Configure MongoDB to make it listen on the controller management IP address. Edit the `/etc/mongodb.conf` file and modify the `bind_ip` key:

```
bind_ip = 10.0.0.11
```

5. Restart the MongoDB service to apply the configuration change:

```
# service mongodb restart
```

6. 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" ]})'
```

7. Configure the Telemetry service to use the database:

Edit the `/etc/ceilometer/ceilometer.conf` file and change the `[database]` section:

```
[database]
# The SQLAlchemy connection string used to connect to the
# database (string value)
connection = mongodb://ceilometer:CEILOMETER_DBPASS@controller:27017/
ceilometer
```

8. 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:

```
# openssl rand -hex 10
```

Edit the `/etc/ceilometer/ceilometer.conf` file and change the `[publisher]` section. Replace `CEILOMETER_TOKEN` with the results of the `openssl` command:

```
[publisher]
# Secret value for signing metering messages (string value)
metering_secret = CEILOMETER_TOKEN
```

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

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

```
[DEFAULT]
...
auth_strategy = keystone
```

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

Edit the `/etc/ceilometer/ceilometer.conf` file and change the `[keystone_auth_token]` section:

```
[keystone_authtoken]
auth_host = controller
auth_port = 35357
auth_protocol = http
auth_uri = http://controller:5000
admin_tenant_name = service
admin_user = ceilometer
admin_password = CEILOMETER_PASS
```

Also set the `[service_credentials]` section:

```
[service_credentials]
os_auth_url = http://controller:5000/v2.0
os_username = ceilometer
os_tenant_name = service
os_password = CEILOMETER_PASS
```

11. Restart the services with their new settings:

```
# service ceilometer-agent-central restart
# service ceilometer-agent-notification restart
# service ceilometer-api restart
# service ceilometer-collector restart
# service ceilometer-alarm-evaluator restart
# service ceilometer-alarm-notifier restart
```

## Install the Compute agent for Telemetry

Telemetry provides an API service that provides a collector and a range of disparate agents. This procedure details how to install the agent that runs on the compute node.

1. Install the Telemetry service on the compute node:

```
# apt-get install ceilometer-agent-compute
```

2. Edit the `/etc/nova/nova.conf` file and add the following lines to the `[DEFAULT]` section:

```
[DEFAULT]
...
instance_usage_audit = True
instance_usage_audit_period = hour
notify_on_state_change = vm_and_task_state
notification_driver = nova.openstack.common.notifier.rpc_notifier
notification_driver = ceilometer.compute.nova_notifier
```

3. Restart the Compute service:

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

4. You must set the secret key that you defined previously. The Telemetry service nodes share this key as a shared secret:

Edit the `/etc/ceilometer/ceilometer.conf` file and change these lines in the `[publisher]` section. Replace `CEILOMETER_TOKEN` with the ceilometer token that you created previously:

```
[publisher]
# Secret value for signing metering messages (string value)
metering_secret = CEILOMETER_TOKEN
```

5. Add the Identity service credentials:

Edit the `/etc/ceilometer/ceilometer.conf` file and change the `[keystone_authtoken]` section:

```
[keystone_authtoken]
auth_host = controller
auth_port = 35357
auth_protocol = http
admin_tenant_name = service
admin_user = ceilometer
admin_password = CEILOMETER_PASS
```

Also set the `[service_credentials]` section:

```
[service_credentials]
os_auth_url = http://controller:5000/v2.0
os_username = ceilometer
os_tenant_name = service
os_password = CEILOMETER_PASS
```

6. Restart the service with its new settings:

```
# service ceilometer-agent-compute restart
```

## Configure the Image Service for Telemetry

1. To retrieve image samples, you must configure the Image Service to send notifications to the bus.

Edit `/etc/glance/glance-api.conf` and modify the `[DEFAULT]` section:

```
notification_driver = messaging
rpc_backend = rabbit
rabbit_host = controller
rabbit_password = RABBIT_PASS
```

2. Restart the Image Services with their new settings:

```
# service glance-registry restart
# service glance-api 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.

Edit `/etc/cinder/cinder.conf` and add in the `[DEFAULT]` section on the controller and volume nodes:

```
control_exchange = cinder
notification_driver = cinder.openstack.common.notifier.rpc_notifier
```





## Next steps

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



# 14. Add the Database service

## Table of Contents

Database service overview .....	105
Install the Database service .....	106
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:

```
# apt-get install python-trove python-troveclient python-glanceclient \
trove-common trove-api trove-taskmanager
```

2. Respond to the prompts for [database management](#) and [\[keystone\\_authtoken\] settings](#), and [API endpoint](#) registration. The `trove-manage db_sync` command runs automatically.

3. Prepare OpenStack:

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

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

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

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





# 15. 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 6, “Configure the Image Service” \[43\]](#) 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
```

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



```
$ neutron net-list
+-----+-----+
+-----+
| id                               | name      | subnets
+-----+-----+
| 3c612b5a-d1db-498a-babb-a4c50e344cb1 | demo-net  | 20bcd3fd-5785-41fe-
ac42-55ff884e3180 192.168.1.0/24
| 9bce64a3-a963-4c05-bfcd-161f708042d1 | ext-net   | b54a8d85-b434-4e85-
a8aa-74873841a90d 203.0.113.0/24
+-----+-----+
+-----+
```

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

- 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](#).

- 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 | vFW7Bp8PQGNo
| config_drive
| created | 2014-04-09T19:24:27Z
| flavor | m1.tiny (1)
| hostId
| id
05682b91-81a1-464c-8f40-8b3da7ee92c5
| 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 | 7cf50047f8df4824bc76c2fdf66d11ec
| updated | 2014-04-09T19:24:27Z
| user_id | 0e47686e72114d7182f7569d70c519c9
+-----+
+-----+

```

## 6. Check the status of your instance:

```

$ nova list
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| ID | Name | Status | Task |
+-----+-----+-----+-----+-----+
| 05682b91-81a1-464c-8f40-8b3da7ee92c5 | demo-instance1 | ACTIVE | -
| Running | demo-net=192.168.1.3 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+

```

The status changes from BUILD to ACTIVE when your instance finishes the build process.



## To access your instance using a virtual console

- Obtain a *Virtual Network Computing (VNC)* session URL for your instance and access it from a web browser:

```
$ nova get-vnc-console demo-instance1 novnc
+-----+
+-----+
+
+ Type | Url
+-----+
+
+ novnc | http://controller:6080/vnc_auto.html?token=2f6dd985-f906-4bfc-
b566-e87ce656375b
+-----+
+-----+
+
```



### Note

If your web browser runs on a host that cannot resolve the *controller* host name, you can replace *controller* with the IP address of the management interface on your controller node.

The CirrOS image includes conventional username/password authentication and provides these credentials at the login prompt. After logging into CirrOS, we recommend that you verify network connectivity using **ping**.

Verify the `demo-net` tenant network gateway:

```
$ ping -c 4 192.168.1.1
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp_req=1 ttl=64 time=0.357 ms
64 bytes from 192.168.1.1: icmp_req=2 ttl=64 time=0.473 ms
64 bytes from 192.168.1.1: icmp_req=3 ttl=64 time=0.504 ms
64 bytes from 192.168.1.1: icmp_req=4 ttl=64 time=0.470 ms

--- 192.168.1.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 2998ms
rtt min/avg/max/mdev = 0.357/0.451/0.504/0.055 ms
```

Verify the `ext-net` external network:

```
$ 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_req=1 ttl=53 time=17.4 ms
64 bytes from 174.143.194.225: icmp_req=2 ttl=53 time=17.5 ms
64 bytes from 174.143.194.225: icmp_req=3 ttl=53 time=17.7 ms
64 bytes from 174.143.194.225: icmp_req=4 ttl=53 time=17.5 ms

--- openstack.org ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3003ms
rtt min/avg/max/mdev = 17.431/17.575/17.734/0.143 ms
```

## To access your instance remotely

1. Add rules to the default security group:

- a. Permit *ICMP* (ping):

```
$ nova secgroup-add-rule default icmp -1 -1 0.0.0.0/0
```

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. Create a *floating IP* address on the `ext-net` external network:

```
$ neutron floatingip-create ext-net
Created a new floatingip:
```

Field	Value
fixed_ip_address	
floating_ip_address	203.0.113.102
floating_network_id	9bce64a3-a963-4c05-bfcd-161f708042d1
id	05e36754-e7f3-46bb-9eaa-3521623b3722
port_id	
router_id	
status	DOWN
tenant_id	7cf50047f8df4824bc76c2fdf66d11ec

3. Associate the floating IP address with your instance:

```
$ nova floating-ip-associate demo-instance1 203.0.113.102
```



### Note

This command provides no output.

4. Check the status of your floating IP address:

```
$ nova list
```

ID	Name	Status	Task
05682b91-81a1-464c-8f40-8b3da7ee92c5	demo-instance1	ACTIVE	-
	Running	demo-net=192.168.1.3, 203.0.113.102	

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

```
$ ping -c 4 203.0.113.102
PING 203.0.113.102 (203.0.113.112) 56(84) bytes of data.
64 bytes from 203.0.113.102: icmp_req=1 ttl=63 time=3.18 ms
64 bytes from 203.0.113.102: icmp_req=2 ttl=63 time=0.981 ms
64 bytes from 203.0.113.102: icmp_req=3 ttl=63 time=1.06 ms
64 bytes from 203.0.113.102: icmp_req=4 ttl=63 time=0.929 ms

--- 203.0.113.102 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
```

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

```
$ ssh cirros@203.0.113.102
The authenticity of host '203.0.113.102 (203.0.113.102)' 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 '203.0.113.102' (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!

## Launch an instance with legacy networking (nova-network)

### 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.

- 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.

- 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](#).

- 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 |
+-----+-----+-----+-----+-----+
| ID | Name | Status | Task |
State | Power State | Networks |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+

```







## Appendix A. Reserved user IDs

In OpenStack, certain user IDs are reserved and used to run specific OpenStack services and own specific OpenStack files. These users are set up according to the distribution packages. The following table gives an overview.



### Note

Some OpenStack packages generate and assign user IDs automatically during package installation. In these cases, the user ID value is not important. The existence of the user ID is what matters.

**Table A.1. Reserved user IDs**

Name	Description	ID
ceilometer	OpenStack Ceilometer Daemons	Assigned during package installation
cinder	OpenStack Cinder Daemons	Assigned during package installation
glance	OpenStack Glance Daemons	Assigned during package installation
heat	OpenStack Heat Daemons	Assigned during package installation
keystone	OpenStack Keystone Daemons	Assigned during package installation
neutron	OpenStack Neutron Daemons	Assigned during package installation
nova	OpenStack Nova Daemons	Assigned during package installation
swift	OpenStack Swift Daemons	Assigned during package installation
trove	OpenStack Trove Daemons	Assigned during package installation

Each user belongs to a user group with the same name as the user.

# 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](#)

- [Virtual Machine Image Guide](#)

The following books explain how to use the OpenStack dashboard and command-line clients:

- [API Quick Start](#)
- [End User Guide](#)
- [Admin User Guide](#)
- [Command-Line Interface Reference](#)

The following documentation provides reference and guidance information for the OpenStack APIs:

- [OpenStack API Complete Reference \(HTML\)](#)
- [API Complete Reference \(PDF\)](#)
- [OpenStack Block Storage Service API v2 Reference](#)
- [OpenStack Compute API v2 and Extensions Reference](#)
- [OpenStack Identity Service API v2.0 Reference](#)
- [OpenStack Image Service API v2 Reference](#)
- [OpenStack Networking API v2.0 Reference](#)
- [OpenStack Object Storage API v1 Reference](#)

The [Training Guides](#) offer software training for cloud administration and management.

## ask.openstack.org

During the set up or testing of OpenStack, you might have questions about how a specific task is completed or be in a situation where a feature does not work correctly. Use the [ask.openstack.org](http://ask.openstack.org) site to ask questions and get answers. When you visit the <http://ask.openstack.org> site, scan the recently asked questions to see whether your question has already been answered. If not, ask a new question. Be sure to give a clear, concise summary in the title and 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.

## OpenStack mailing lists

A great way to get answers and insights is to post your question or problematic scenario to the OpenStack mailing list. You can learn from and help others who might have similar issues. To subscribe or view the archives, go to <http://lists.openstack.org/cgi-bin/mailman/listinfo/openstack>. You might be interested in the other mailing lists for specific projects or development, which you can find [on the wiki](#). A description of all mailing lists is available at <http://wiki.openstack.org/MailingLists>.

## 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 \(trove\)](#)

- [Bugs: Orchestration \(heat\)](#)
- [Bugs: Telemetry \(ceilometer\)](#)
- [Bugs: Queue Service \(marconi\)](#)
- [Bugs: OpenStack API Documentation \(api.openstack.org\)](#)
- [Bugs: OpenStack Documentation \(docs.openstack.org\)](#)

## The OpenStack IRC channel

The OpenStack community lives in the #openstack IRC channel on the Freenode network. You can hang out, ask questions, or get immediate feedback for urgent and pressing issues. To install an IRC client or use a browser-based client, go to <http://webchat.freenode.net/>. You can also use Colloquy (Mac OS X, <http://colloquy.info/>), mIRC (Windows, <http://www.mirc.com/>), or XChat (Linux). When you are in the IRC channel and want to share code or command output, the generally accepted method is to use a Paste Bin. The OpenStack project has one at <http://paste.openstack.org>. Just paste your longer amounts of text or logs in the web form and you get a URL that you can paste into the channel. The OpenStack IRC channel is #openstack on irc.freenode.net. You can find a list of all OpenStack IRC channels at <https://wiki.openstack.org/wiki/IRC>.

## Documentation feedback

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](#).

## OpenStack distribution packages

The following Linux distributions provide community-supported packages for OpenStack:

- **Debian:** <http://wiki.debian.org/OpenStack>
- **CentOS, Fedora, and Red Hat Enterprise Linux:** <http://openstack.redhat.com/>
- **openSUSE and SUSE Linux Enterprise Server:** <http://en.opensuse.org/Portal:OpenStack>
- **Ubuntu:** <https://wiki.ubuntu.com/ServerTeam/CloudArchive>

# Glossary

## API

Application programming interface.

## 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.

## external network

A network segment typically used for instance Internet access.

## firewall

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.

#### Icehouse

Project name for the ninth release of OpenStack.

#### 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.

Object Storage

The OpenStack core project that provides eventually consistent and redundant storage and retrieval of fixed digital content. The project name of OpenStack Object Storage is swift.



## OpenStack

OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a data center, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface. OpenStack is an open source project licensed under the Apache License 2.0.

## Orchestration

An integrated project that orchestrates multiple cloud applications for OpenStack. The project name of Orchestration is heat.

## path MTU discovery (PMTUD)

Mechanism in IP networks to detect end-to-end MTU and adjust packet size accordingly.

## plug-in

Software component providing the actual implementation for Networking APIs, or for Compute APIs, depending on the context.

## promiscuous mode

Causes the network interface to pass all traffic it receives to the host rather than passing only the frames addressed to it.

## public key authentication

Authentication method that uses keys rather than passwords.

## RESTful

A kind of web service API that uses REST, or Representational State Transfer. REST is the style of architecture for hypermedia systems that is used for the World Wide Web.

## router

A physical or virtual network device that passes network traffic between different networks.

## security group

A set of network traffic filtering rules that are applied to a Compute instance.

## service catalog

Alternative term for the Identity Service catalog.

## subnet

Logical subdivision of an IP network.

## Telemetry

An integrated project that provides metering and measuring facilities for OpenStack. The project name of Telemetry is ceilometer.

## tenant

A group of users, used to isolate access to Compute resources. An alternative term for a project.

## trove

OpenStack project that provides database services to applications.

## 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.

