



# Reference Architecture for OpenStack Folsom

---

OpenStack is one of the fastest growing open source projects today, with thousands of active developers and hundreds of actively supporting companies and individuals. Getting started with OpenStack has been simplified by the existence of this large development community to such a degree that a single script can turn a virtual machine into a usable OpenStack-based test environment. But in attempting to make a production-ready system, the choices suddenly narrow and deploying a system that can be managed after production starts is even more difficult.

The Cisco Reference Architecture for OpenStack Folsom is one of the current models for getting to that manageable deployment. It includes a model for the compute, network, and storage components, the virtualization platform, and also how to install the system. A set of guidelines and a validated process for using well-known Hypervisor and OpenStack code distributions, along with a tested and validated DevOps deployment model based on the work of the community (driven by Puppet Labs with their Puppet DevOps toolset), is included. This document will also cover the tests that have been completed against the environment to highlight where the system is functional, and in some cases, areas of possible concern where improvements will be made.

## Reference Architecture

The reference architecture is intended to support a common model for OpenStack deployments, meeting an average set of user and usage requirements. Aspects of the system can be modified and scaled either up or down, depending on specific end user needs, but the system is based on a few key tenets:

1. Middle of the system specs to meet “average” scenarios.
2. Provide for future extension to support a high availability service (not covered in this document).
3. Minimize system differences to reduce Day0 (initial) and Day 2 (break/fix) operational requirements.
4. Avoid top bin devices/components to manage overall system cost and improve system value.
5. Cover the principal OpenStack core components.
6. Enhance the core components when appropriate (for services such as furuter deployments in a High Availability or added functionality).

While the reference architecture is intended to be able to support all core services in the current OpenStack project, minor gaps exist in support between the core project and the current tested reference. Specifically, due to Cisco’s commitment to efficiently delivering value to its customers, no testing / validation has occurred for Swift or Cinder projects and support for the Nova network elements has been

deprecated for Quantum services only. Testing included in this document demonstrates that although Cinder has not been tested as part of Cisco OpenStack Edition, installation and deployment of Cinder is still possible atop the reference base install.

## Reference Architecture Target Requirements

In addition to the general model requirements, we will also address a set of feature/capability requirements in future testing iterations using Cisco Reference Architecture for OpenStack Folsom. Namely:

- Ability to support highly available access to OpenStack services (Nova, Glance, Keystone, Quantum, Swift, and Cinder).
- Support for micro, small, medium, large, x-large, and xx-large instances (up to 32GB memory and 8 cores per VM).
- Support for local instance persistence.
- Ability to support VM migration and restart.
- Support for single interface or multi-interface networks.
- Support for bonded physical network interface resilience.
- Support OpenStack Essex and Folsom releases against Ubuntu 12.04 LTS and RedHat RHEL 6.

As noted above, the model and our validation diverge in areas where we do not yet have support from the Cisco Reference Architecture testing and validation efforts.

## Physical Infrastructure Model

To simplify operational management, only two types of systems are included in the model: compute-centric ([Table 1](#)) and storage-centric ([Table 2](#)).

**Table 1** Compute Model based on UCS C220-M3

Element	Type	Quantity	Description
CPU	Intel E5-2660	2	Mid-tier high core count CPU for a balance of power and VM scale.
Memory	1600MHz 16GB dual rank DIMM	16	Supports up to 4 xx-large instances per physical system.
NIC	Cisco VIC	1	Provides dual port 10G interfaces for resiliency and the ability to support VM-FEX for hypervisor bypass (on supported Hypervisors).
Disk Controller	Mega-RAID 9266i	1	Provide memory-backed RAID operation across the local disks, improve performance of lower cost/denser disk options. RAID 10 for performance
Disk Drives	600GB 10Krpm SAS	8	Provide a large possible footprint for local VM instances with reasonable performance.

The compute system is based on the 1RU C220-M3 platform and leverages a low power 8 core CPU and 256GB of memory giving a memory-to-core ratio of 16:1. The storage subsystem is based on a high performance RAID controller and 8 SAS disks for a flexible model for distributed CINDER and/or Ceph

storage. The network interface is based on the Cisco Virtual Interface Controller (VIC), providing dual 10Gbps network channels and enabling Hypervisor Bypass with Virtual Machine Fabric Extension (VM-FEX) functionality when combined with a Nexus 5500 series data center switch as the Top of Rack (TOR) device, Fibre Channel over Ethernet (FCOE) storage, and Network Interface Card (NIC) bonding for network path resiliency or increased network performance for video streaming, high performance data moves, or storage applications.

**Table 2** *Storage Model based on UCS C240-M3*

Element	Type	Quantity	Description
CPU	Intel E5-2609	2	Lower core count CPU for a reduced computational non-complex workload.
Memory	1600MHz 8GB dual rank DIMM	4	Provides working memory for in-system disk cache.
NIC	Cisco VIC	1	Provides dual port 10G interfaces for bonded NIC resiliency.
Disk Controller	Mega-RAID 9266i	1	Provide memory-backed RAID operation across the local disks for non-Swift-based storage. No RAID config.
Disk Drives	1 TB 7.2Krpm SATA-6	24	Disk for Swift or block or NAS depending on usage model.

The storage system is based on the 2RU C240-M3 platform, which is similar at the baseboard level to the C220-M3, but provides up to 24 2.5" drive slots. With 24 spindles, this platform is focused on storage as opposed to compute, and while it could be used as configured for a combined all-in-one platform, the reference makes use of dual low power 4 core CPUs, and a much smaller memory space at 32GB total, which is our current model for managing SWIFT or CINDER-focused nodes specifically. This platform also includes the Cisco VIC for up to 20Gbps of storage forwarding with link resiliency when combined with the dual TOR model.

## Compute BIOS

The current default host BIOS configuration is appropriate for deployment; however, it is convenient to change some of the parameters to accelerate boot, and address hardware differences (such as having the Cisco FlexFlash installed). This will improve the automated deployment process. The manual steps required to implement these changes can be found in the Cisco UCS C-Series Servers Integrated Management Controller CLI Configuration Guide or the Cisco UCS C-Series Servers Integrated Management Controller Configuration Guide for CLI or Web UI based configuration.

Some of the non-standard recommended parameters are defined in [Table 3](#):

**Table 3** *Non-Standard Compute BIOS Recommended Parameters*

BIOS Parameter	Value	Description
bios/LomOpromControlPort0	Disabled	Disable un-used LOM port BIOS. There are either 2 or 4 of these (0,1 or 0,1,2,3) for C220 vs. C240.
bios/UsbPortInt	Disabled	Access to the internal 4GB USB card if installed (not in the reference model).
bios/UsbPortSdCard	Disabled	Access to the internal 16GB Cisco FlexFlash if installed (not in the reference model).

**Table 3** *Non-Standard Compute BIOS Recommended Parameters (continued)*

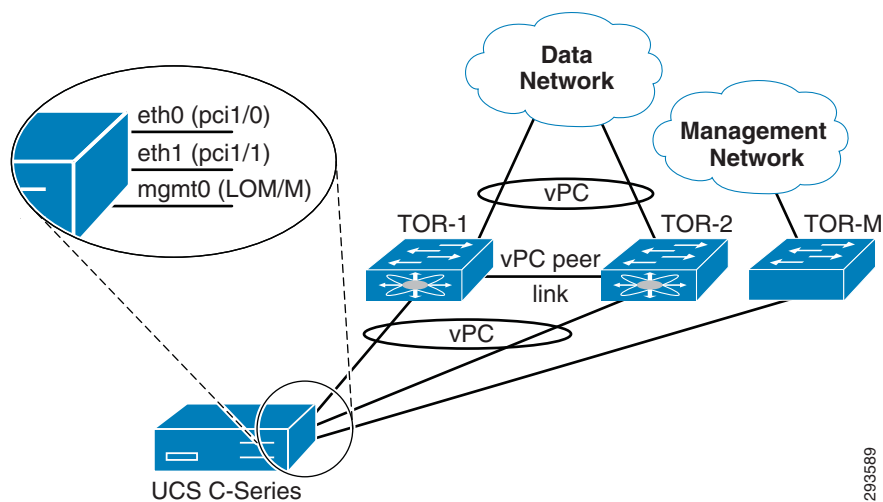
Boot-order	cdrom,pxe,hdd	Simplifies day 0 OS install, PXE being the principal method for installing the hypervisor OS.
BIOS/CIMC Version	1.5(f)	The current 1.5 BIOS release provides enhanced CIMC-based access to BIOS and PCI device management.

A set of utility scripts are available to facilitate BIOS updates and storage configurations, along with collecting data needed for the reference deployment model. These scripts are available from the Cisco Systems GitHub repository:

<https://github.com/CiscoSystems/ucs-openstack-cimc-expect.git>

## Network Model

The upstream network is based on the Nexus 5500 series switch enabling the use of a number of advanced scale-out network services in the Layer 2 (Link Local) and Layer 3 (Routed Network) services. In the basic reference, the TOR switches are configured as a virtual Port Channel (vPC) pair, with a set of 10Gigabit connections between them as the vPC peer link, and a set of two ports each as the vPC to the upstream device (a Nexus 7000 series device in our Virtual Multiservice Datacenter (VMDC) network model). [Figure 1](#) shows the normal link configurations, and [Table A-1](#) in [Appendix A](#) shows an example port mapping for the two TOR switches with the server management interfaces (TOR-M in [Figure 1](#)) collapsed onto the TOR-1 and TOR-2 devices.

**Figure 1** *Physical Network Configuration*

The physical model above makes the following assumptions:

- VLAN 100 is for management, both CIMC (power, KVM, and bare metal management) and for the hypervisor management interface.
- An additional VLAN (or in certain network models, additional VLANs) will be provisioned on the trunk ports by the Cisco Quantum plugin.
- vPC is used to connect to an upstream router (or routers) who support LACP port channels at a minimum.

- vPC is used to connect to the physical servers, but this is either manually or automatically configured after the base OS is installed due to the inability of most Linux operating systems to PXE boot and configure a bonded network interface.

Logically, the network is segregated either via VLANs (as proposed in this reference) or via an overlay L2 technology like VXLAN. The latter is expected to become the standard mechanism for providing much greater tenancy or project-level L2 segregation scale by the Fall 2013 OpenStack release time frame (the Havana release). L3 and other network services have only limited support in the current Folsom and Grizzly releases, and in our Folsom reference architecture, the L3 services will be supported by the L3\_Agent network model. Security will be provided by the IPtables security and NAT functionality driven by Quantum.

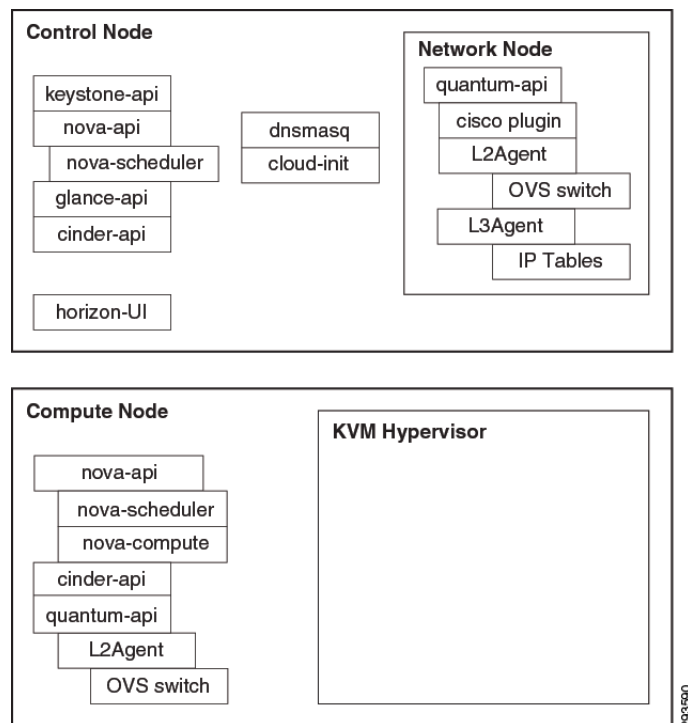
The reference architecture will work with any vPC enabled TOR switch, but we recommend either the Nexus 3548 switch with L3 services if performance is the principal concern, or the Nexus 5548-UP with the L3 daughter card if features like VM-FEX and/or FCOE capability are of interest. A BOM for both devices is included in Appendix B.

## Software Architecture

The system software architecture for the Folsom release of COE on Ubuntu 12.04 LTS is straight forward (Figure 2). Future references will be released to support the Grizzly release, and RedHat-based Hypervisor OS systems.

The non-HA model has a single node acting as a control node, running the control APIs, with the exception of nova-compute. This same system also runs the Quantum L2 and L3 agents providing local to external network connectivity and security.

**Figure 2 Primary Software Components on Control and Compute Nodes**



The current deployment model also includes a build node, which provides a number of additional services beyond the OpenStack-specific deployment components highlighted above, namely:

- Cobbler (<https://github.com/cobbler>) to support bare metal install of the compute node hypervisor and control node base OS.
- Puppet Agent and Puppet Master (<https://puppetlabs.com>) DevOps toolset for system management and deployment services.
- Nagios, Collectd, and Graphite for system health data collection from the control and compute nodes.

The control and compute nodes also run a Puppet Agent process (not highlighted above) that periodically synchronizes with the Puppet Master to enable central system level changes.

An update to this design document will cover the deployment and capabilities of a high availability model control plane and will add in the advanced storage services (SWIFT and an external CINDER service, specifically).

## Small Scale Example System

As most deployments don't immediately jump to a full rack or even multi-rack OpenStack systems deployment, we have provided a small system model that should allow a simple update to a complete rack model. This simple model starts with a single switch and a set of three compute class nodes.

The following system is an example of a small test or application support deployment setup model. This system is designed to support ~48 virtual compute nodes across this particular environment and to be configured into a single 42 RU server rack with the intention of having additional devices added to build out a full server rack-based system (Table 4).

**Table 4** *Small Scale System Rack Layout*

Location in Rack (RU number)	Principal Function	Component Name	Element
Slot 42 (top of rack)	Network	TOR-1	5548-UP
Slot 41-39	Expansion		Blank
Slot 38	Control/Network	build-server	C220-M3 Compute
Slot 37	Control/Network	control-server	Blank
Slot 35-36	Expansion		Blank
Slot 34	Compute	build-server 01	C220-M3 Compute
Slot 33	Compute	build-server 02	C220-M3 Compute

The physical network model for this system is also quite simple, but is designed to be ready to scale out to the full rack system as well and uses the same port model, but without any of the redundant links installed initially (and no virtual port channel or redundant uplinks configured). Refer to Appendix A for the wiring diagram.

# Rack Scale Example System

This is a system that includes all of the components (Table 5) needed for the upcoming HA systems model, along with the SWIFT and CINDER storage components. It should be able to provide on the order of ~480 virtual machines and has both scalable block and object storage architectures.

**Table 5** *Rack Scale Example System*

Location in Rack (RU number)	Principal Function	Component Name	Element
Slot 42 (top of rack)	Network	TOR-1	5548-UP
Slot 41	Network	TOR-2	5548-UP
Slot 40			Expansion
Slot 39			Expansion
Slot 38	Build	build-server	C220-M3 Compute
Slot 37	Control/Compute	control-server01	C220-M3 Compute
Slot 36	Control/Compute	control-server02	C220-M3 Compute
Slot 35	Control/Compute	control-server03	C220-M3 Compute
Slot 34	Compute	compute-server01	C220-M3 Compute
Slot 33	Compute	compute-server02	C220-M3 Compute
Slot 32	Compute	compute-server03	C220-M3 Compute
Slot 31	Compute	compute-server04	C220-M3 Compute
Slot 30	Compute	compute-server05	C220-M3 Compute
Slot 29	Compute	compute-server06	C220-M3 Compute
Slot 28	Compute	compute-server07	C220-M3 Compute
Slot 27	Compute	compute-server08	C220-M3 Compute
Slot 26	Compute	compute-server09	C220-M3 Compute
Slot 25	Compute	compute-server10	C220-M3 Compute
Slot 24	Compute	compute-server11	C220-M3 Compute
Slot 23	Compute	compute-server12	C220-M3 Compute
Slot 22	Compute	compute-server13	C220-M3 Compute
Slot 21	Compute	compute-server14	C220-M3 Compute
Slot 20	Compute	compute-server15	C220-M3 Compute
Slot 19	Storage Proxy	proxy-server01	C220-M3 Compute
Slot 18	Storage Proxy	proxy-server02	C220-M3 Compute
Slot 17	Storage Proxy	proxy-server03	C220-M3 Compute
Slot 15-16	Cinder Block	block-server01	C240-M3 Compute
Slot 13-14	Cinder Block	block-server02	C240-M3 Compute
Slot 11-12	Swift	swift-server01	C240-M3 Compute
Slot 9-10	Swift	swift-server02	C240-M3 Compute
Slot 7-8	Swift	swift-server03	C240-M3 Compute

**Table 5** Rack Scale Example System (continued)

Location in Rack (RU number)	Principal Function	Component Name	Element
Slot 5-6	Swift	swift-server04	C240-M3 Compute
Slot 3-4	Swift	swift-server05	C240-M3 Compute
Slot 1-2			Expansion

## Systems Installation

The following section walks through the software steps required to install an OpenStack system. This process presumes an environment as described above.

### Assumptions

Although other configurations are supported, the following instructions target an environment with a build node, a controller node, and at least one compute node. Additional compute nodes may optionally be added.

Also, these instructions primarily target deployment of OpenStack onto UCS servers (either blades or rack-mount form factors). Several steps in the automation leverage the UCS manager or CIMC to execute system tasks. Deployment on non-UCS gear may well work, particularly if the gear has functional IPMI, but may require additional configuration or additional manual steps to manage systems.

COE Folsom requires that you have two physically or logically (VLAN) separated IP networks. You must have an external router or L3 switch that provides connectivity between these two networks. This is required in order to support Quantum and the Cloud-Init metadata server for provider router-based deployments (similar to FlatDHCP + Floating IP support in the nova-network model):

[http://docs.openstack.org/folsom/openstack-network/admin/content/adv\\_cfg\\_l3\\_agent\\_metadata.html](http://docs.openstack.org/folsom/openstack-network/admin/content/adv_cfg_l3_agent_metadata.html)

One network is used to provide connectivity for OpenStack API endpoints, Open vSwitch (OVS) GRE endpoints and OpenStack/UCS management. The second network is used by OVS as the physical bridge interface and by Quantum as the public network.

## Creating a Build Server

To create a build server, perform the following:

- 
- Step 1** To deploy Cisco OpenStack, first configure a build server.
- This server has relatively modest hardware requirements: 2 GB RAM, 20 GB storage, Internet connectivity, and a network interface on the same network as the eventual management interfaces of the OpenStack cluster machines are the minimal requirements. This machine can be physical or virtual; eventually a pre-built VM of this server will be provided, but this is not yet available.
- Step 2** Install Ubuntu 12.04 LTS onto this build server.
- A minimal install with openssh-server is sufficient. Configure the network interface on the OpenStack cluster management segment with a static IP.



Also, when partitioning the storage, choose a partitioning scheme which provides at least 15 GB free space under /var, as installation packages and ISO images used to deploy OpenStack will eventually be cached there.

**Step 3** When the installation finishes, log in and become root: **sudo -H bash**



**Note** If you have proxies, or your control and compute nodes do not have Internet access, note the following:

- If you require a proxy server to access the Internet, be aware that proxy users have occasionally reported problems during the phases of the installation process that download and install software packages. A common symptom of proxy trouble is that apt will complain about hash mismatches or file corruptions when verifying downloaded files. A few known scenarios and workarounds include:
  - If the apt-get process reports a "HASH mismatch," you may be facing an issue with a caching engine. If it's possible to do so, bypassing the caching engine may resolve the problem.
- If you do have a proxy, at a minimum, you will want to export the two types of proxies needed in your root shell when running fetch commands, as noted in the relevant sections.
- You will also want to change the \$proxy setting in site.pp to reflect your local proxy.
- Another possible change is if you don't have "public" Internet accessible IPs for all of your machines (build, control, compute, etc.) and are building this in a controlled environment. If this is the case, ensure that \$default\_gateway is *\*not\** set in site.pp and all of the files required for installing the control and compute nodes will be fetched from the boot server.
- You have two choices for setting up the build server. You can follow the manual steps below, or you can run a one line script that tries to automate this process. In either case, you should end up with the Puppet modules installed, and a set of template site manifests in /etc/puppet/manifests.

**Step 4 Model 1:** Run the Script.

- a. To run the install script, copy and paste the following on your command line (as root with your proxy set if necessary, as above):

```
curl -s -k -B
https://raw.githubusercontent.com/CiscoSystems/folsom-manifests/multi-node/install_os_puppet |
/bin/bash
```

#### With a Proxy, Use

```
https_proxy=http://proxy.example.com:80/ curl -s -k -B
https://raw.githubusercontent.com/CiscoSystems/folsom-manifests/multi-node/install_os_puppet >
install_os_puppet
chmod +x install_os_puppet
./install_os_puppet -p http://proxy.example.com:80/
```

- b. Jump to "[Customizing the Build Server, page -10](#)" or continue with [Step 5](#).

**Step 5 Model 2:** Run the Commands Manually.

All should now install any pending security updates:

```
apt-get update && apt-get dist-upgrade -y && apt-get install -y puppet git ipmitool
debmirror
```




---

**Note** The system may need to be restarted after applying the updates.

---

- a. Get the reference example manifests. Under the [folsom-manifests GitHub repository](#), you will find different branches, so select the one that matches your topology plans most closely. In the following examples, the multi-node branch will be used, which is likely the most common topology:

```
git clone https://github.com/CiscoSystems/folsom-manifests ~/cisco-folsom-manifests/
cd ~/cisco-folsom-manifests
git checkout -q 2012.2.3
```

#### With a Proxy

```
https_proxy=http://proxy.example.com:80 git clone
https://github.com/CiscoSystems/folsom-manifests ~/cisco-folsom-manifests/
cd ~/cisco-folsom-manifests
https_proxy=http://proxy.example.com:80 git checkout -q 2012.2.3
```

- b. Copy the Puppet manifests from ~/cisco-folsom-manifests/manifests/ to /etc/puppet/manifests/

```
cp ~/cisco-folsom-manifests/manifests/* /etc/puppet/manifests
```

- c. Then get the reference Puppet modules from Cisco's GitHub repository:

```
(cd /etc/puppet/manifests; sh /etc/puppet/manifests/puppet-modules.sh)
```

#### With a Proxy

```
(cd /etc/puppet/manifests; http_proxy=http://proxy.example.com:80
https_proxy=http://proxy.example.com:80 sh /etc/puppet/manifests/puppet-modules.sh)
```

---

## Customizing the Build Server

In the /etc/puppet/manifests directory you will find these files:

- clean\_node.sh
- cobbler-node.pp
- core.pp
- modules.list
- puppet-modules.sh
- reset\_nodes.sh
- site.pp.example

At a high level, cobbler-node.pp manages the deployment of cobbler to support booting of additional servers into your environment. The core.pp manifest defines the core definitions for OpenStack service deployment. The site.pp.example manifest captures the user-modifiable components and defines the various parameters that must be set to configure the OpenStack cluster, including the Puppet Master and Cobbler setup on the build server. clean\_node.sh is a shell script provided as a convenience to deployment users; it wraps several cobbler and Puppet commands for ease of use when building and rebuilding the nodes of the OpenStack cluster. reset\_nodes.sh is a wrapper around clean\_node.sh to rebuild your entire cluster quickly with one command.

**Note**

IMPORTANT! You must copy `site.pp.example` to `site.pp` and then edit it as appropriate for your installation. It is internally documented.

```
cp /etc/puppet/manifests/site.pp.example /etc/puppet/manifests/site.pp
vi /etc/puppet/manifests/site.pp
```

Then, use the ‘puppet apply’ command to activate the manifest:

```
puppet apply -v /etc/puppet/manifests/site.pp
```

When the ‘puppet apply’ command runs, the Puppet client on the build server will follow the instructions in the `site.pp` and `cobbler-node.pp` manifests and will configure several programs on the build server:

- **Network Time Protocol daemon (NTPD)**—a time synchronization server used on all OpenStack cluster nodes to ensure time throughout the cluster is correct.
- **tftpd-hpa**—a TFTP server used as part of the PXE boot process when OpenStack nodes boot up
- **dnsmasq**—a DNS and DHCP server used as part of the PXE boot process when OpenStack nodes boot up.
- **Cobbler**—an installation and boot management daemon, which manages the installation and booting of OpenStack nodes.
- **apt-cacher-ng**—a caching proxy for package installations, used to speed up package installation on the OpenStack nodes.
- **Nagios**—an infrastructure monitoring application, used to monitor the servers and processes of the OpenStack cluster.
- **Collectd**—a statistics collection application, used to gather performance and other metrics from the components of the OpenStack cluster.
- **Graphite and Carbon**—a real-time graphing system for parsing and displaying metrics and statistics about OpenStack.
- **Apache**—a web server hosting sites to implement Graphite, Nagios, and Puppet web services.

The initial Puppet configuration of the build server will take several minutes to complete as it downloads, installs, and configures all the software needed for these applications.

Once the `site.pp` manifest has been applied to your system, you need to stage Puppet plugins so they can be accessed by the managed nodes:

```
puppet plugin download
```

After the build server is configured, the systems listed in `site.pp` should be defined in cobbler on the build server:

```
cobbler system list
control-server
compute-server01
compute-server02
```

And now, you should be able to use cobbler to build your controller:

```
/etc/puppet/manifests/clean_node.sh {node_name}
```

**Note**

Replace `node_name` with the name of your controller.

`clean_node.sh` is a script which does several things:

- Configures Cobbler to PXE boot the specified node with appropriate PXE options to do an automated install of Ubuntu.
- Uses Cobbler to power-cycle the node.
- Removes any existing client registrations for the node from Puppet, so Puppet will treat it as a new install.
- Removes any existing key entries for the node from the SSH known hosts database.

You can watch the progress on the console of your controller node as Cobbler completes the automated install of Ubuntu. Once the installation finishes, the controller node will reboot and then will run Puppet after it boots up. Puppet will pull and apply the controller node configuration defined in the Puppet manifests on the build server.

This step will take several minutes, as Puppet downloads, installs, and configures the various OpenStack components and support applications needed on the control node. `/var/log/syslog` on the controller node will display the progress of the Puppet configuration run.



#### Note

It may take more than one Puppet run for the controller node to be set up completely, especially if there are proxies in the path as some proxies can have issues with apt-get installs and updates. Observe the log files (`/var/log/syslog` on the controller node) to verify that the controller configuration has converged completely to the configuration defined in Puppet.

Once the Puppet configuration of the controller has completed, follow the same steps to build each of the other nodes in the cluster, using `clean_node.sh` to initiate each install. As with the controller, the other nodes will take several minutes for Puppet configuration to complete and may require multiple runs of Puppet before they are fully converged to their defined configuration state.

As a short cut, if you want to build all of the nodes defined in your `cobbler-node.pp` file, you can run:

```
for n in `cobbler system list`; do clean_node.sh $n ; done
```

Or you can run a full reset script, which also does this and re-runs the build-node Puppet apply and Puppet plugin download steps:

```
./reset_nodes.sh
```

Once the OpenStack nodes have been built using Cobbler, run Puppet on the build node a second time:

```
puppet agent -t
```

This second Puppet run will gather information about the individual OpenStack nodes collected by Puppet when they were being built, and use that information to set up status monitoring of the OpenStack cluster on the build server.

## Testing OpenStack

Once the nodes are built, and once Puppet runs have completed on all nodes (watch `/var/log/syslog` on the cobbler node), you should be able to log into the OpenStack Horizon interface:

<http://ip-of-your-control-node> / user: admin, password: Cisco123 (if you didn't change the defaults in the `site.pp` file).

You will still need to log into the console of the control node to load in an image using user: localadmin, password: ubuntu. If you SU to root, you will need to source the `openrc` auth file, which is in the root's home directory (run "source `openrc`" in `/root/`), and you can launch a test file in `/tmp/nova_test.sh`.

## Deploy Your First VM

The following manual deployment procedure should be used after completing clean puppet runs on OpenStack Nodes and restarting quantum-server and quantum-plugin-openvswitch-agent services.

**Step 1** Create quantum public network.

```
quantum net-create public --router:external=True
```

**Step 2** We are using 192.168.221.0/24 as our external network.



**Note** The eth settings on the Controller Node associated to this network should not have an IP address assigned to it as it will function in bridged mode.

```
quantum subnet-create public 192.168.221.0/24
```



**Note**

If there are upstream routers/L3 switches that use HSRP/GLBP/VRRP that use low-order IP addresses such as .2 and .3 then the default subnet address assignments used by Quantum for this subnet (such as floating IP addresses and the Qrouter interface [default is .3]) will directly conflict with these real IP addresses on the upstream first hop routers. You can alter these default address assignments for the Quantum subnet by using the "--allocation-pool" range when creating the Quantum subnet. The example that follows will use the default upstream router address of .1 (in this example the upstream HSRP address would be 192.168.221.1) and the first addresses for floating-IPs will begin at .10:

```
quantum subnet-create --allocation-pool
start=192.168.221.10,end=192.168.221.250 public 192.168.221.0/24
```

**Step 3** Create the internal (data) network used for Tenants. Create additional networks and associated subnets as needed. In the example below, we are assigning specific DNS servers that will be used by the instances.

```
quantum net-create net10
quantum subnet-create net10 10.10.10.0/24 --dns_nameservers list=true 8.8.8.8 8.8.4.4
```



**Note**

Replace 8.8.8.8 8.8.4.4 with the IP address(es) of the DNS server or servers virtual machines should use.

**Step 4** Create a virtual router and an associated interface used for the subnet created in the previous step:

```
quantum router-create router1
quantum router-interface-add router1 <net10-subnet-uuid>
```

**Step 5** Connect the virtual router to your external network:

```
quantum router-gateway-set router1 <public-net-id>
```

You must create a static route on your physical router/layer-3 switch to the subnet that you define for your Tenants (quantum subnet-create command). The next-hop for the static route is the quantum router public interface. Use the following command to find the public address of your Quantum router:

```
quantum port-list -- --device_id <router-id> --device_owner network:router_gateway
```

- a. Continuing with our example, here is the required static route that **MUST** be configured on the external physical router/layer-3 switch. If you used the before mentioned "--allocation-pool" flag then the first address in the pool will be used for the qrouter address. So, in that example the next-hop entry for the example below would be "192.168.221.10":

```
ip route 10.10.10.0 255.255.255.0 192.168.221.3
```

- Step 6** Download an image and add it to Glance:

#### For Ubuntu Precise

```
wget
http://cloud-images.ubuntu.com/precise/current/precise-server-cloudimg-amd64-disk1.img
glance add name="precise-x86_64" is_public=true container_format=ovf disk_format=qcow2
< precise-server-cloudimg-amd64-disk1.img
```

#### For Cirros Cloud Image

```
wget http://download.cirros-cloud.net/0.3.1/cirros-0.3.1-x86_64-disk.img
glance add name="cirros-x86_64" is_public=true disk_format=qcow2 container_format=ovf
< cirros-0.3.1-x86_64-disk.img
```

- Step 7** Create an SSH keypair and add the public key to Nova. Make sure you create a key-pair for your Network and Controller Nodes. Note: leave the passphrase empty when creating the keypair:

```
ssh-keygen
cd /root/.ssh/
nova keypair-add --pub_key id_rsa.pub <key_name>
```

- Step 8** Boot an Instance (Precise image example):

```
quantum net-list
nova boot --image precise-x86_64 --flavor m1.tiny --key_name <key_name> --nic
net-id=<quantum-net10-id> <your_instance_name>
```

#### Cirros Image Example

```
nova boot --image cirros-x86_64 --flavor m1.tiny --key_name <key_name> --nic
net-id=<quantum-net10-id> <your_instance_name>
```

- Step 9** Verify that your instance has spawned successfully:

```
nova show <your_instance_name>
```

- Step 10** Verify connectivity to Instance from the node running Quantum L3 Agent (Controller Node). Since we are using namespaces, we run the commands from the context of the router using the "ip netns exec qrouter" syntax. Below, we list the qrouter to get its router-id, we connect to the qrouter and get a list of its addresses, we ping the instance from the qrouter and then we SSH into the instance from the qrouter:

```
quantum router-list
ip netns exec qrouter-<quantum-router-id> ip addr list
ip netns exec qrouter-<quantum-router-id> ping <fixed-ip-of-instance>
ip netns exec qrouter-<quantum-router-id> ssh ubuntu@<fixed-ip-of-instance>
```



#### Note

You can get the internal fixed IP of your instance with the following command: `nova show <your_instance_name>`

- Step 11** Create and associate a Floating IP. You will need to get a list of the networks copy the correct IDs:

```
quantum net-list
quantum floatingip-create <public/ext-net-id>
quantum floatingip-associate <floatingip-id> <internal VM port-id>
```

```

or in a single step:
quantum net-list
quantum port-list
quantum floatingip-create --port_id <internal VM port-id> <public/ext-net-id>

```

**Step 12** Enable Ping and SSH to Instances:

```

nova secgroup-add-rule default icmp -1 -1 0.0.0.0/0
nova secgroup-add-rule default tcp 22 22 0.0.0.0/0

```

**Step 13** Ping and SSH to your Instances from an external host.

## A Scripted Example

For those that might want to automate the process above, note item 5a. about adding a specific IP route to the upstream router from your public facing network. A scripted example is available for download here:

<https://github.com/CiscoSystems/quantum-l3-test>

And instructions for running the script are included in the Readme file.

# Running your OpenStack Environment

If the previous set of instructions was followed correctly, you should now have a simple system image, security and network access, and manipulate and map storage to running instances. Future OpenStack white papers will look at managing users and projects, scaling systems, running multi-site systems, and a host of other operations and scale out tasks. Check back often for more from the OpenStack team.

## Switch Port Mapping

Table 6 provides an example of switch port mapping.

**Table 6** Example Switch Port Mapping

Switch Port	Trunk/Access	VLANs	Speed	Connected Device	Device Port	Port-Channel	vPC
		native,allowed					
e0/1	trunk	100,all	10G	Upstream	PC-100 p1	98	98
e0/2	trunk	100,all	10G	Upstream	PC-100 p2	98	98
e0/3	trunk	100,all	10G	TOR-2	e0/3	99	99
e0/4	trunk	100,all	10G	TOR-2	e0/4	99	99
e0/5	trunk	100,100	10G	compute-00	pci1/0	100	100
e0/6	trunk	100,100	10G	compute-01	pci1/0	101	101
e0/7	trunk	100,100	10G	compute-02	pci1/0	102	102
e0/8	trunk	100,100	10G	compute-03	pci1/0	103	03
e0/9	trunk	100,100	10G	compute-04	pci1/0	104	104

Table 6 Example Switch Port Mapping (continued)

Switch Port	Trunk/Access	VLANs	Speed	Connected Device	Device Port	Port-Channel	vPC
e0/10	trunk	100,100	10G	compute-05	pci1/0	105	105
e0/11	trunk	100,100	10G	compute-06	pci1/0	106	106
e0/12	trunk	100,100	10G	compute-07	pci1/0	107	107
e0/13	trunk	100,100	10G	compute-08	pci1/0	108	108
e0/14	trunk	100,100	10G	compute-09	pci1/0	109	109
e0/15	trunk	100,100	10G	compute-10	pci1/0	110	110
e0/16	trunk	100,100	10G	compute-11	pci1/0	111	111
e0/17	trunk	100,100	10G	compute-12	pci1/0	112	112
e0/18	trunk	100,100	10G	compute-13	pci1/0	113	113
e0/19	trunk	100,100	10G	compute-14	pci1/0	114	114
e0/20	trunk	100,100	10G	compute-15	pci1/0	115	115
e0/21	trunk	100,100	10G	compute-16	pci1/0	116	116
e0/22	trunk	100,100	10G	compute-17	pci1/0	117	117
e0/23	trunk	100,100	10G	compute-18	pci1/0	118	118
e0/24	trunk	100,100	10G	compute-19	pci1/0	119	119
e0/25	trunk	100,100	10G	storage-00	pci1/0	120	120
e0/26	trunk	100,100	10G	storage-01	pci1/0	121	121
e0/27	trunk	100,100	10G	storage-02	pci1/0	122	122
e0/28	trunk	100,100	10G	storage-03	pci1/0	123	123
e0/29	trunk	100,100	10G	storage-04	pci1/0	124	124
e0/30	trunk	100,100	10G	storage-05	pci1/0	125	125
e0/31	trunk	100,100	10G	storage-06	pci1/0	126	126
e0/32	trunk	100	1G	compute-00	lot-m		
e0/33	access	100	1G	compute-02	lot-m		
e0/34	access	100	1G	compute-04	lot-m		
e0/35	access	100	1G	compute-06	lot-m		
e0/36	access	100	1G	compute-08	lot-m		
e0/37	access	100	1G	compute-10	lot-m		
e0/38	access	100	1G	compute-12	lot-m		
e0/39	access	100	1G	compute-14	lom-m		
e0/40	access	100	1G	compute-16	lom-m		
e0/41	access	100	1G	compute-18	lom-m		
e0/42	access	100	1G	storage-00	lom-m		
e0/43	access	100	1G	storage-01	lom-m		
e0/44	access	100	1G	storage-02	lom-m		
e0/45	access	100	1G	storage-03	lom-m		



**Table 6** Example Switch Port Mapping (continued)

Switch Port	Trunk/Access	VLANs	Speed	Connected Device	Device Port	Port-Channel	vPC
e0/46	access	100	1G	storage-04	lom-m		
e0/47	access	100	1G	storage-05	lom-m		
e0/48	access	100	1G	storage-06	lom-m		

## Bills of Material

Table 7, Table 8, Table 9, and Table 10 lists the following bills of material for consideration.

- “Compute Reference Bill of Materials”
- “Storage Reference Bill of Materials”
- “Network TOR Model A (Nexus 3548) Reference Bill of Materials”
- “Network TOR Model B (Nexus 5548-UP) Reference Bill of Materials”

**Table 7** Compute Reference Bill of Materials

Product	Description	Quantity
UCSC-C220-M3S	UCS C220 M3 SFF w/o CPU, mem, HDD, PCIe, PSU, w/ rail kit	1
UCS-CPU-E5-2660	2.20 GHz E5-2660/95W 8C/20MB Cache/DDR3 1600MHz	2
UCS-MR-1X162RY-A	16GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	16
A03-D600GA2	600GB 6Gb SAS 10K RPM SFF HDD/hot plug/drive sled mounted	8
UCS-RAID-9266-NB	MegaRAID 9266-8i with no battery back up	1
R2XX-RAID10	Enable RAID 10 Setting	1
UCSC-PCIE-C10T-02	Cisco VIC 1225T Dual Port 10GBaseT CAN	1
UCSC-PSU-650W	650W power supply for C-series rack servers	2
CAB-C13-C14-2M	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	2
UCSC-DLOM-01	Dedicated LOM Mode BIOS setting for C-Series Servers	1
UCSC-HS-C220M3	Heat Sink for UCS C220 M3 Rack Server	2
UCSC-RAIL1	Rail Kit for C220, C22, C24 rack servers	1

**Table 8** Storage Reference Bill of Materials

Product	Description	Quantity
UCSC-C240-M3S	UCS C240 M3 SFF w/o CPU, mem, HD, PCIe, w/ rail kit, expdr	1
UCS-CPU-E5-2609	2.4 GHz E5-2609/80W 4C/10MB Cache/DDR3 1066MHz	2
UCS-MR-1X082RY-A	8GB DDR3-1600-MHz RDIMM/PC3-12800/dual rank/1.35v	4
UCS-HDD300GI2F105	300GB 6Gb SAS 15K RPM SFF HDD/hot plug/drive sled mounted	24
UCS-RAID-9266-NB	MegaRAID 9266-8i with no battery back up	1
UCSC-PCIE-C10T-02	Cisco VIC 1225T Dual Port 10GBaseT CNA	1

**Table 8 Storage Reference Bill of Materials (continued)**

Product	Description	Quantity
UCSC-PSU-650W	650W power supply for C-series rack servers	2
CAB-C13-C14-2M	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	2
UCSC-DLOM-01	Dedicated LOM Mode BIOS setting for C-Series Servers	1
UCSC-HS-C240M3	Heat Sink for UCS C240 M3 Rack Server	2
UCSC-RAIL-2U	2U Rail Kit for UCS C-Series servers	1
UCSC-PCIF-01F	Full height PCIe filler for C-Series	3

**Table 9 Network TOR Model A (Nexus 3548) Reference Bill of Materials**

Product	Description	Quantity
N3K-C3548P-FA-L3A	Nexus 3548+L3+Algo, Fwd Airflow (port side exhaust), AC	1
CAB-C13-C14-2M	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	2
GLC-T	1000BASE-T SFP	16
SFP-H10GB-CU1M	10GBASE-CU SFP+ Cable 1 Meter	16
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	16
N3KUK9-503A1.1	NX-OS Release 5.0(3)A1(1)	1
DCNM-LAN-N3K-K9	DCNM for LAN Advanced Edt. for Nexus 3000	1
N2200-PAC-400W	N2K/N3K AC Power Supply, Std airflow (port side exhaust)	2
N3548-ALGK9	Nexus 3548 Algo Boost License	1
N3548-BAS1K9	Nexus 3548 Base License	1
N3548-LAN1K9	Nexus 3548 Layer 3 LAN Enterprise License	1
N3K-C3064-ACC-KIT	Nexus 3064PQ Accessory Kit	1
NXA-FAN-30CFM-F	Nexus 2K/3K Single Fan, std airflow (port side exhaust)	4

**Table 10 Network TOR Model B (Nexus 5548-UP) Reference Bill of Materials**

Product	Description	Quantity
N5K-C5548UP-FA	Nexus 5548 UP Chassis, 32 10GbE Ports, 2 PS, 2 Fans	1
N5548P-FAN	Nexus 5548P Fan Module	2
N55-PAC-750W	Nexus 5500 PS, 750W, Front to Back Airflow	2
CAB-C13-C14-2M	Power Cord Jumper, C13-C14 Connectors, 2 Meter Length	2
GLC-T	1000BASE-T SFP	8
SFP-H10GB-CU1M	10GBASE-CU SFP+ Cable 1 Meter	16
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	8
N55-D160L3-V2	Nexus 5548 Layer 3 Daughter Card, Version 2	1
N55-M16UP	Nexus 5500 Unified Mod 16p 10GE Eth/FCoE OR 16p 8/4/2/1G FC	1
GLC-T	1000BASE-T SFP	8

**Table 10** Network TOR Model B (Nexus 5548-UP) Reference Bill of Materials (continued)

Product	Description	Quantity
SFP-H10GB-CU3M	10GBASE-CU SFP+ Cable 3 Meter	8
N5KUK9-602N1.2	Nexus 5000 Base OS Software Rel 6.0(2)N1(2)	1
N55-LAN1K9	Layer 3 License for Nexus 5500 Platform	1
N55-BAS1K9	Layer 3 Base License for Nexus 5500 Platform	1
N5548-ACC-KIT	Nexus 5548 Chassis Accessory Kit	1

## Tested Design

The design tested is a subset of the current rack level reference architecture, principally due to the current gaps between overall functionality and supportable tested efforts in the COE distribution. The test system was built out of a set of five UCS C-200 servers, connected together per [Figure 1](#). The 10 Gbps links provide the primary data networks, while the 1 Gbps links provide a management plane including access to the Cisco Integrated Management Controller for each node.

## OpenStack Test Cases

The testing described in this paper demonstrates that OpenStack is a flexible Cloud operating system that can be deployed quickly and in an automated fashion. Cisco OpenStack Edition powered by Puppet provides a flexible deployment model for OpenStack that allows cloud administrators to deploy OpenStack networking, compute, storage, and identity services on multiple compute nodes with minimal effort. To highlight functional use cases, the following tests were performed:

1. [OpenStack Installation in Multiple Deployment Models, page -19](#)
2. [OpenStack Installation using Multiple Network Models, page -26](#)
3. [OpenStack Installation with Cinder Block Storage Service, page -28](#)
4. [OpenStack Installation Automated with Minimal Input, page -29](#)
5. [Monitoring Functional State with an Administrative Dashboard, page -31](#)
6. [OpenStack Manual Installation Minimum Deployment Time, page -33](#)
7. [OpenStack Application Infrastructure Deployment Minimum Deployment Time, page -33](#)

## OpenStack Installation in Multiple Deployment Models

OpenStack installation in multiple deployment models (multi-node and single node/all-in-one) use case is the deployment by a system administrator of an installation type.

Deploy OpenStack in a number of different ways, including:

1. Standalone on a single physical server.
  - [Standalone Test Setup, page -20](#)
  - [Standalone Test Procedure, page -21](#)
  - [Pass/Fail Criteria, page -26](#)
2. Multi-Node fashion on physical compute, network, and storage elements.

- [Multi-Node Test Setup, page -22](#)
- [Multi-Node Test Procedure, page -23](#)
- [Multi-Node Test Case Troubleshooting, page -25](#)
- [Pass/Fail Criteria, page -26](#)

## Standalone Test Setup

The topology for this test includes a total of 5 UCS C200M2 servers connected to a Nexus 3064 top-of-rack switch pair. The first C200 will serve as a “build” node to serve DHCP/PXE and Puppet for installing the other servers. The second C200 will serve as the control node on which coordinating services for the cloud will be run (including Horizon, the OpenStack GUI). The remaining C200s serve as compute nodes on which instances may be launched.

Each C200 is equipped with two LoM (LAN on motherboard) gigabit Ethernet (GbE) interfaces and two 10-gigabit Ethernet interfaces. The first GbE interface (“eth0”) is cabled to a management network. The C200 Cisco Integrated Management Controller (CIMC) is set to use this adapter as its network interface rather than as a dedicated management interface. The remaining gigabit Ethernet interface is not used in this test. Each of the two 10-gigabit Ethernet interfaces is cabled to a separate Nexus 3064 top-of-rack (ToR) switch. All server ports on the ToR switches should be configured in VLAN 1:

```
interface Ethernet1/1
  description cvf2-server-f1 10Gb-1
  switchport mode trunk
  switchport trunk allowed vlan 1
  spanning-tree port type edge trunk
  flowcontrol receive on
```

The VLAN interface on each ToR switch should carry a separate IP address:

```
cvf2-leaf-f1# show run int Vlan1

!Command: show running-config interface Vlan1
!Time: Sun Mar 24 13:24:44 2013

version 5.0(3)U5(1a)

interface Vlan1
  ip address 2.6.1.1/24
  ip dhcp relay address 2.63.1.33
  no shutdown

cvf2-leaf-f1#
cvf2-leaf-f2# show run int Vlan1

!Command: show running-config interface Vlan1
!Time: Sun Mar 24 13:24:52 2013

version 5.0(3)U5(1a)

interface Vlan1
  ip address 2.6.129.1/24
  ip dhcp relay address 2.63.1.33
  no shutdown

cvf2-leaf-f2#
```

## Standalone Test Procedure

The following procedure for standalone OpenStack deployment on a single physical server was performed:

- 
- Step 1** Install Ubuntu 12.04 Server on the first node in the cluster (e.g., the “build node”).
- Step 2** As root, run the following commands to apply updates and install a few necessary software packages:
- export http\_proxy=http://proxy-rtp-1.cisco.com:8080
  - export https\_proxy=https://proxy-rtp-1.cisco.com:8080
  - apt-get update
  - apt-get dist-upgrade -y
  - apt-get install git ipmitool debmirror
- Step 3** Fetch the Cisco OpenStack Edition – Folsom Puppet modules by running the following commands as root:
- git clone -b multi-node <https://github.com/CiscoSystems/folsom-manifests> ~/cisco-folsom-manifests/
  - cp -R ~/cisco-folsom-manifests/\* /etc/puppet/manifests/
  - cd /etc/puppet/manifests
  - sh puppet-modules.sh
- Step 4** Copy the site.pp.example file to site.pp and edit it to customize it for your environment. You will supply information about the remaining nodes in your cloud, networking information, and other configuration information.
- cp site.pp.example site.pp
  - vi site.pp
- Step 5** If necessary, edit cobbler-node.pp to add any late commands you would like to use to customize your bare-metal operating system installs.
- Step 6** Remove the proxy variables from your environment and apply the edited site.pp file. Once the site.pp file has been applied, make puppet plugins available for nodes to download:
- unset http\_proxy; unset https\_proxy
  - puppet apply -v site.pp
  - puppet plugin download
- Step 7** Install the control node by issuing the following commands to power cycle it. When the machine reboots, it should begin a PXE installation of Ubuntu 12.04. You may wish to launch the KVM console from the CIMC before issuing these commands so you can observe the process as it takes place.
- cobbler system poweroff controller
  - cobbler system poweron controller
- Step 8** When the control node has completed the operating system install, log in to it as localadmin with password “ubuntu.” If you disabled puppet from running at startup via your site.pp file and need to make any final changes before running the installation process, do so now.
- Step 9** If you disabled puppet from running at startup via your site.pp file, edit /etc/default/puppet and change “no” to “yes,” then run the following command as root:
- service puppet start

- Step 10** Observe the puppet agent logging messages in `/var/log/syslog` as it installs and configures the OpenStack control node. You should not see any error messages.
- Step 11** Once the control node has completed its installation process, boot the compute nodes in the cluster. As they boot, they will begin PXE installation of Ubuntu 12.04. You may wish to access the KVM consoles of each system prior to issuing these commands to observe the install process.
- `cobbler system poweroff compute-01; cobbler system poweron compute-01`
  - `cobbler system poweroff compute-02; cobbler system poweron compute-02`
  - `cobbler system poweroff compute-03; cobbler system poweron compute-03`
- Step 12** Once the operating install process completes, log in to each node as `localadmin` with password “ubuntu.” If you chose to disable puppet from running at startup via your `site.pp` file and need to make any final changes to each node before starting the OpenStack installation process, do so now.
- Step 13** If you chose to disable puppet from running at startup via your `site.pp` file, start the puppet agent by running the following command as root on each node:
- `service puppet start`
- Step 14** Observe the puppet log messages in `/var/log/syslog` on each system. You should not see any errors.
- Step 15** Once the compute nodes have finished their automated setup process, run the following command on the control node and verify that each compute node is listed in the output:
- `nova-manage host list`
- Step 16** Access Horizon by opening a web browser and typing in the IP address or hostname of your control node. Verify that you can log in as “admin” with password “Cisco123.”
- Step 17** Download a test script to your control node and execute it with the following commands:
- `source ~/openrc`
  - `https_proxy=https://proxy-rtp-1.cisco.com:8080 git clone https://github.com/CiscoSystems/quantum-l3-test`
  - `cd quantum-l3-test`
  - `./create_vm`
- Step 18** Verify that you can now log in to the instance created by the test script via SSH to either its public or private IP address.
- Step 19** Verify that you access the Nagios monitoring display by opening a web browser and typing “[http://\[hostname or IP address of your build node\]/nagios3](http://[hostname or IP address of your build node]/nagios3)” into the URL bar. When prompted, log in as “admin” with password “Cisco123.” Verify that you can see basic server health parameters under the “Services” section in Nagios.
- Step 20** Power off one of the compute nodes. Verify that the “down” status is reflected in Nagios.
- Step 21** Boot the compute node back up. Verify that its state changes to “up” in Nagios.
- 

## Multi-Node Test Setup

The topology for this test consists of a single C200 M2 server. The server is connected to the network by a single GbE LAN on Motherboard (LoM) interface.

## Multi-Node Test Procedure

The following procedure for multi-node OpenStack deployment on physical compute, network, and storage elements was performed.

- 
- Step 1** Begin by installing Ubuntu 12.04 Server on the C200 M2.
- Step 2** Use the commands below to install updates and necessary software packages:
- a. `export http_proxy=http://proxy-rtp-1.cisco.com:8080`
  - a. `export https_proxy=https://proxy-rtp-1.cisco.com:8080`
  - b. `apt-get update`
  - c. `apt-get dist-upgrade -y`
  - d. `apt-get install puppet screen git -y`
- Step 3** Fetch the puppet code that will automate the bulk of the deployment.
- a. `git clone https://github.com/robertstarmar/folsome-manifests -b all-in-one all-in-one-manifests`
  - b. `cp all-in-one-manifests/* /etc/puppet/manifests/`
  - c. `cd /etc/puppet/manifests/`
  - d. `./puppet-modules.sh`
- Step 4** Remove some of the baseline modules and install patched all-in-one modules.
- a. for n in nova quantum openstack horizon;do; `rm -rf /etc/puppet/modules/$n; git clone -b folsom https://github.com/robertstarmar/puppet-$n /etc/puppet/modules/$n`; done
- Step 5** Copy the example site.pp file and edit it to make adjustments specific to your environment.
- a. `cp /etc/puppet/manifests/site.pp.aio.example /etc/puppet/manifests/site.pp`
  - b. `vi /etc/puppet/manifests/site.pp`
- Step 6** As root, modify the network setup to move the IP address on eth0 to the external bridge that will be created by the puppet modules.
- a. `vi /etc/network/interfaces`

The completed file should look similar to this:

```
# The loopback network interface
auto lo
iface lo inet loopback
# The primary network interface
auto eth0
iface eth0 inet static
    address 0.0.0.0
auto br-ex
iface br-ex inet dhcp
```

- Step 7** Modify `/etc/init/failsafe.conf`
- Step 8** to ensure that the system boots more quickly.

```
root@cvf2-server-f1:/etc/puppet/manifests# grep -A3 'Waiting for network'
/etc/init/failsafe.conf
    $PLYMOUTH message --text="Waiting for network configuration..." || :

    $PLYMOUTH message --text="Waiting up to 60 more seconds for network
configuration..." || :
```

```

$PLYMOUTH message --text="Booting system without full network
configuration..." || :
root@cvf2-server-f1:/etc/puppet/manifests#

```

**Step 9** Run the `site.pp` manifest in “noop” mode to check for any major problems. A few errors here related to packages that must be installed are expected since in “noop” mode the packages won’t actually be installed.

- a. `cd /etc/puppet/manifests`
- b. `puppet apply -v --noop site.pp`

**Step 10** Apply the `site.pp` manifest from a screen session. Observe puppet log messages to verify that there are no major errors. You may notice an error of the form “err: Failed to apply catalog: No title provided and “Host[]” is not a valid resource reference.” This error can be ignored for now.

- a. `screen`
- b. `cd /etc/puppet/manifests`
- c. `puppet apply -v /etc/puppet/manifests/site.pp`

**Step 11** Edit `/etc/hosts`. If you noticed a single line at the end of the file consisting entirely of one IP address, modify the line to include a valid hostname. If definitions for localhost are missing, add them. The completed file should look similar to this:

```

localhost 127.0.1.1 cvf2-server-f1.cisco.com cvf2-server-f1

# The following lines are desirable for IPv6 capable hosts ::1
ip6-localhost ip6-loopback fe00::0 ip6-localnet ff00::0
ip6-mcastprefix ff02::1 ip6-allnodes ff02::2 ip6-allrouters
192.168.2.119 cvf2-server-f1 cvf2-server-f1.cisco.com

```

**Step 12** Save a copy of `/etc/hosts` to root’s home directory and reboot.

- a. `cp /etc/hosts /root/hosts`
- b. `reboot`

**Step 13** When the operating system has booted up again, reapply the puppet manifest and restart networking by running the following command as root:

- a. `puppet apply -v /etc/puppet/manifests/site.pp`
- b. `/etc/init.d/networking stop`
- c. `/etc/init.d/networking start`

**Step 14** Verify that the `br-ex` interface has an IP address by running the following command:

- a. `Ifconfig br-ex`

**Step 15** Copy the hosts file you saved earlier back into place by running the following command as root:

- a. `cp /root/hosts /etc/hosts`

**Step 16** Log in to Horizon by opening a web browser and typing the IP address or hostname of your server into the URL bar. When prompted, use “admin” as the user name and “Cisco123” as the password.

**Step 17** Download and run a test script to create a Quantum router, public and private networks, and launch an instance connected to them.

- a. `export https_proxy=https://proxy-rtp-1.cisco.com:8080`
- b. `export http_proxy=http://proxy-rtp-1.cisco.com:8080`
- c. `git clone https://github.com/CiscoSystems/quantum-l3-test`
- d. `cd quantum-l3-test`



- e. `./create_vm`
- Step 18** Verify that you can SSH to the instance created by the test script via its public or private IP address.
- Step 19** Verify that you can access Nagios by opening a web browser and typing “[http://\[hostname or IP address of your server\]/nagios3](http://[hostname or IP address of your server]/nagios3)” into the URL bar. When prompted, log in as “admin” with password “Cisco123.”
- Step 20** Verify that you can see basic server health parameters under the “Services” section in Nagios.

## Multi-Node Test Case Troubleshooting

A few known issues exist with the installation process of this writing: the most notable being that `/etc/hosts` gets corrupted by Puppet due to a bug in the all-in-one puppet manifests. These manifests have not yet been released by Cisco and this issue is expected to be resolved prior to their official release. The corrupted `/etc/hosts` file will also result in a Puppet error of the form:

```
err: Failed to apply catalog: No title provided and "Host[]" is not a valid
resource reference
```

Additionally, the user may see some Puppet errors pertaining to database setup. These are likely related to a timing issue with various parts of the pre-release Puppet code and will generally be resolved by the running the “puppet apply” command a second time. The errors in question include:

```
Stage[main]/Glance::Db::Mysql/Mysql::Db[glance]/Database[glance]/ensure:
change from absent to present failed: Execution of '/usr/bin/mysql
-NBe create database `glance` character set latin1' returned 1:
ERROR 1045 (28000): Access denied for user 'root'@'localhost'
(using password: NO)

err: /Stage[main]/Nova::Db::Mysql/Mysql::Db[nova]/Database[nova]/ensure:
change from absent to present failed: Execution of '/usr/bin/mysql
-NBe create database `nova` character set latin1' returned 1:
ERROR 1045 (28000): Access denied for user 'root'@'localhost'
(using password: NO)

err: /Stage[main]/Quantum::Db::Mysql/Mysql::Db[quantum]/Database[quantum]/ensure:
change from absent to present failed: Execution of '/usr/bin/mysql
-NBe create database `quantum` character set latin1' returned 1:
ERROR 1045 (28000): Access denied for user 'root'@'localhost'
(using password: NO)

err: /Stage[main]/Keystone::Db::Mysql/Mysql::Db[keystone]/Database[keystone]/ensure:
change from absent to present failed: Execution of '/usr/bin/mysql
-NBe create database `keystone` character set latin1' returned 1:
ERROR 1045 (28000): Access denied for user 'root'@'localhost'
(using password: NO)

err: /Stage[main]/Keystone/Exec[keystone-manage db_sync]: Failed to
call refresh: keystone-manage db_sync returned 1 instead of one
of [0] at /etc/puppet/modules/keystone/manifests/init.pp:161
err: /Stage[main]/Glance::Registry/Exec[glance-manage db_sync]:
Failed to call refresh: glance-manage version_control 0 &&
glance-manage db_sync returned 1 instead of one of [0] at
/etc/puppet/modules/glance/manifests/registry.pp:49
```

## Pass/Fail Criteria

The following pass/Fail criteria was used.

1. Puppet should run without errors that prevent OpenStack from functioning. For 1.2, some errors described in this document may be visible, but should not prevent all software components from being deployed if the procedures in this document are followed.
2. User should be able to log in to Horizon and see instances once booted.
3. User should be able to see accurate up/down status in Nagios.
4. User should be able to SSH in to instances once booted.
5. User should be able to successfully create Quantum routers, networks, subnets, and floating IP addresses.

## OpenStack Installation using Multiple Network Models

OpenStack installation using multiple network models use case is the network deployment by a system administrator of an installation type.

Deploy OpenStack in a number of different ways, including:

1. Multiple physical network interfaces for management and tenant or other networks with GRE tunnels used for isolation on one or more virtual or physical OpenStack elements.
  - [Multiple Physical Network Interfaces Test Setup, page -26](#)
  - [Multiple Physical Network Interfaces Procedure, page -26](#)
  - [Pass/Fail Criteria, page -27](#)
2. Single physical network interface for management and tenant or other networks with VLANs used for tenant isolation on a virtual or physical OpenStack element.
  - [Single Physical Network Interfaces Test Setup, page -27](#)
  - [Single Physical Network Interfaces Test Procedure, page -27](#)
  - [Pass/Fail Criteria, page -27](#)

### Multiple Physical Network Interfaces Test Setup

Multiple physical network interfaces for management and tenant or other networks with GRE tunnels used for isolation on one or more virtual or physical OpenStack elements is documented at:

<http://docwiki.cisco.com/wiki/OpenStack:Folsom-Multinode>

In this setup we use a 5-node topology as described in [OpenStack Installation in Multiple Deployment Models, page -19](#).

### Multiple Physical Network Interfaces Procedure

The following procedure for multiple physical network interfaces OpenStack deployment was performed:

Testing requires at least 3 servers, each with two interfaces, connected by a ToR switch.

- 
- Step 1** Install a minimal install of Ubuntu 12.04 LTS on the first server (the “build” server).
  - Step 2** Add the git and puppet and ipmitool packages to the build server.

- Step 3** `cd /root; git clone -b multi-node https://github.com/CiscoSystems/folsom-manifests; cd folsom-manifests ; git checkout 2012.2.3 ; cp manifests/* /etc/puppet/manifests`
- Step 4** `cd /etc/puppet/manifests.`
- Step 5** Edit site.pp as appropriate for your install.
- Step 6** `sh puppet-modules.sh.`
- Step 7** Run puppet to configure your build node: `puppet apply -v -d /etc/puppet/manifests/site.pp.`
- Step 8** Puppet plugin download.
- Step 9** `cobbler system list` to verify.
- Step 10** `./clean_node.sh` controller to install the controller.
- Step 11** After the controller installs and reboots, log into it and monitor the automatic puppet run that sets it up.
- Step 12** Once it finishes, repeat steps 10-11 for the compute nodes.
- Step 13** Run `puppet agent -t -d` on the build node to finish configuring Nagios.
- 

### Single Physical Network Interfaces Test Setup

This setup varies from a [Multiple Physical Network Interfaces Test Setup](#) only in that it requires a single ethernet interface (“eth0”) with VLAN tags used to separate traffic instead of multiple physical interfaces.

### Single Physical Network Interfaces Test Procedure

Start with the same basic configuration as multiple physical network interfaces, but with these variations.

- Each machine needs one network interface and the interfaces are attached to trunk ports in the ToR switch.
- Two VLANs are passed to each server. One, the “management” VLAN, will be used for internal traffic while the other, the “public” VLAN, will be used for access from the larger “external” network outside of the ToR switch.
  - The “public” VLAN needs to be routable from whatever “external” network should be able to access the VMs created in the OpenStack cluster.
  - The “management” VLAN does not need to be reachable outside of the ToR switch.

Repeat the steps as in the [Multiple Physical Network Interfaces Procedure](#), but using VLAN interfaces.

### Pass/Fail Criteria

To test both multiple and single physical network test scenarios, verify that the following is possible once configuration of the nodes has finished:

1. Add precise and cirros images in Glance.
2. Configure a quantum router with two networks, one internal to the VMs and one public for floatingip access to the VMs.
3. Boot at least two VMs, with customized properties assigned either via `config_drive` or via metadata service.
4. Access Horizon to display and use the console of each VM.

5. Verify ping access from the namespace of the quantum router to each VM.
6. Verify ssh access from the namespace of the quantum router to each VM.
7. Verify ssh access between VMs in their private network.
8. Adjust default security policies to allow SSH and ICMP to VMs.
9. Assign a floatingip to each VM.
10. Verify access to each VM via its “public” floatingip from the external network.
11. Attempt to launch 18 more VMs and validate that quota and resource restrictions are enforced, preventing all 20 VMs from being launched.

## OpenStack Installation with Cinder Block Storage Service

OpenStack installation with Cinder block storage service use case is the storage deployment by a system administrator of an installation type.

A Cisco OpenStack Edition user should be able to configure Cinder to use a logical volume on the control server in which to place cinder volumes.

### Cinder Block Storage Service Test Setup

OpenStack installation with Cinder block storage service assumes a working Cisco OpenStack Edition 2012.2.3 installation with at least one VM created and running. Although not officially released as part of Cisco OpenStack Edition, Cisco has made available an experimental Cinder Puppet module in version 2012.2.3.

To add actual usable storage, a storage point needs to be identified and made available as a logical volume for cinder to use. In this test case, a 20GB loopback filesystem on a loopback volume called “cinder-volumes” is created on the cloud controller node. In this test case, a 20GB test volume is created in /var/tmp.

### Cinder Block Storage Service Procedure

The following procedure for OpenStack installation with cinder block storage service was performed.



#### Note

All commands in the test procedure should be run as root.

- Step 1** Create the empty loopback file.
 

```
dd if=/dev/zero of=cinder.loopback bs=1 count=0 seek=20G
```
- Step 2** Create a loopback device using the created loopback file.
 

```
losetup /dev/loop2 cinder.loopback
```
- Step 3** Create a physical volume based on the loopback device.
 

```
pvcreate /dev/loop2
```
- Step 4** Create a volume group named "cinder-production" from the physical volume.
 

```
vgcreate cinder-production /dev/loop2
```
- Step 5** Edit /etc/cinder/cinder.conf.
  - a. change volume\_group=cinder-volumes to volume\_group=cinder-production

- Step 6** Restart the cinder services.
- service cinder-volume restart
  - service cinder-api restart
  - service cinder-scheduler restart
- Step 7** Create a 1GB cinder volume:
- ```
cinder create --display_name myVolume 1
```
- Step 8** View the created volume.
- ```
cinder list
```
- Step 9** Once the volume is successfully created (status: available), attach it to a VM as device vdc.
- ```
nova volume-attach <vmname> <cinder volume ID> /dev/vdc
```
- Step 10** Ensure the volume is attached.
- ```
cinder list (status: attached)
```
- Step 11** Test the storage attached to the VM.
- SSH to the VM.
  - Check that vdc exists.
    - cat /proc/partitions
  - Make a filesystem on vdc and mount.
    - mkdir /my\_storage
    - mkfs.ext3 /dev/vdc
    - mount /dev/vdc /my\_storage
  - Confirm that the volume is usable.
    - touch /my\_storage/testfile
    - echo "foo" > /my\_storage/testfile
    - cat /my\_storage/testfile (should print foo to the console)
  - Unmount the volume.

```
umount /my_storage
```
- Step 12** Exit SSH to the VM.
- Step 13** Unattach the storage from the VM and delete the volume.
- ```
nova volume-detach <myname> <cinder volume ID>  
cinder delete <cinder volume ID>
```
- 

### Pass/Fail Criteria

The test is successful if all steps complete without error.

## OpenStack Installation Automated with Minimal Input

OpenStack installation automated with minimal input use case is deployment automation by a system administrator of an installation type.

The installation runs automatically once basic information is provided to the installer that does not exceed the following:

1. Network address information.
2. Physical device address and access information (e.g., power management, MAC address).
3. User names and passwords if necessary for external or API access.
4. Deployment model information (network interface bonding, disk use models, etc.).
5. Network and storage control interfaces for external devices.
6. Default management information for automatically deplorable network and storage elements.

### Automated Installation Test Setup

Refer to [OpenStack Installation in Multiple Deployment Models, page -19](#). This test case requires only the prerequisites for whichever OpenStack installation method you choose.

### Automated Installation Procedure

The following procedure for OpenStack installation automated with minimal input was performed.

Follow the instructions in [OpenStack Installation in Multiple Deployment Models, page -19](#) to obtain a copy of `site.pp.example`.

All user-configurable install instructions reside in `site.pp.example`, which after editing, must then be copied to `site.pp` for execution.

`site.pp.example` is well-documented internally. The general information that needs to be known to correctly modify the file is as follows:

- Whether or not you need to use a proxy to reach the Internet. For convenience, Cisco offers the options of both HTTP and FTP for package retrieval.
- An accessible NTP server.
- The general network configuration you wish to use.
- The IP addresses, hostnames, and domain names you wish to use for your build, control, and compute nodes.
- The physical hardware in each node that will connect to the respective physical/logical network (e.g., network interfaces).
- The MAC addresses for each node.
- The netmask and default gateway used for the private network.
- IP and login information for power control of the physical nodes.

By default, Cisco pre-populates most of this information. Cisco also defines one of each node: build, controller, and compute. To add additional nodes, you must add additional node information as described in `site.pp.example`.

Complete the installation process described in [OpenStack Installation in Multiple Deployment Models, page -19](#).

### Pass/Fail Criteria

The test case passes if the installation procedure in [Standalone Test Procedure of OpenStack Installation in Multiple Deployment Models, page -19](#) can be completed without adding additional information to `site.pp` beyond the node information described in the procedure.

## Monitoring Functional State with an Administrative Dashboard

OpenStack monitoring functional state with an administrative dashboard use case is a management and monitoring feature by a system administrator showing system functional state.

The functional state of the system should be available to a systems state dashboard of some nature allowing for a view of areas of function or areas where a system maintainable or repair task is required. Failures should be reported in a time commensurate with their severity. Capacity limitations and system bottlenecks should also be discoverable via these mechanisms.

It should be possible to see:

### Control

1. Total server down alerts

### Compute

1. Total server down alerts

### Network

1. Total server down alerts
2. Rabbit messaging alerts are working
3. Rabbit API alerts are working

### System

1. Total server down alerts
2. Keystone service not running alerts
3. Keystone authentication not working alerts
4. Keystone endpoints not defined alerts
5. NTP times not sufficiently synchronized alerts
6. nova-api service not running alerts
7. nova-api service not answering queries alerts
8. glance-api not running alerts
9. glance-api misconfigured alerts
10. glance alerts if no images found
11. apache for Horizon not running triggers alerts
12. VM status logs cpu / memory consumption are gathered

## Test Setup

Testing should be done with a standard multi-node test pod. The pod will be installed using puppet in the standard COE multinode deployment as documented at:  
<http://docwiki.cisco.com/wiki/OpenStack:Folsom-Multinode>

To prepare for testing, configure puppet with information about the build node, the control node, and two compute nodes. Then, deploy only the build node.

## Functional State with Administrative Dashboard Procedure

Perform the following monitoring functional state with an administrative dashboard procedure.

### Control

---

- Step 1** Deploy the control node. After the puppet run finishes, power off the node and verify Nagios alerting about control node.
- 

### Compute

---

- Step 1** Deploy the compute node. After the puppet run finishes, power off the node and verify Nagios alerting about compute node.
- 

### Network

---

- Step 1** When executing 5.0.1, verify that network connectivity Nagios alerts are triggered.
- Step 2** Configure an incorrect `rabbit_virtual_host` in `/etc/nova/nova.conf` and restart nova services on the controller. Verify “RabbitMQ Alive” Nagios alerts triggered.
- Step 3** Configure an incorrect `rabbit_password` in `/etc/nova/nova.conf` and restart nova services on the controller. Verify “RabbitMQ Objects” Nagios alerts triggered.
- 

### System

---

- Step 1** When executing 5.0.1, verify that system tests (load, users, disk space) generate Nagios alerts.
- Step 2** Stop the keystone service on the control node and verify the Keystone alert is triggered.
- Step 3** Configure an incorrect password in `/etc/nrpe.d/check_keystone.cfg` on the control node and verify the Keystone alert is triggered.
- Step 4** Delete endpoints for services from Keystone and verify the Keystone alert is triggered.
- Step 5** Stop NTP services throughout the cluster and change a system’s clock by 10 minutes. Confirm that Nagios NTP alerts are triggered by the clock drift.
- Step 6** Stop the nova-api service on the controller and verify that the Nova-Api alert is triggered.
- Step 7** Configure an incorrect password in `/etc/nrpe.d/check_novaapi.cfg` on the control node and verify the Nova-Api alert is triggered.
- Step 8** Stop the glance-api service on the controller and verify that the Glance alert is triggered.
- Step 9** Configure an incorrect password in `/etc/nrpe.d/check_glance.cfg` on the control node and verify the Glance alert is triggered.
- Step 10** Confirm that the Glance alert logs errors until an image called “precise-x86\_64” has been uploaded.
- Step 11** Stop the apache2 service on the control node and verify the HTTP alert is triggered.



**Step 12** Launch a couple of VMs and verify that CPU / memory statistics are logged under VM Stats alerts.

---

### Pass/Fail Criteria

The test case passes if all steps above succeed.

## OpenStack Manual Installation Minimum Deployment Time

OpenStack manual installation minimum deployment time use case is a systems deployment times measure by a system administrator showing deployment times.

The system should deploy rapidly. The following are target times that should be possible to meet:

- Initial setup of a node from scratch, assuming all addressing and other manual entry information exists: 120 minutes
- Additional node setup: 60 minutes
- Network setup: 60 minutes
- Additional network setup: 60 minutes
- Storage environment setup: 120 minutes
- Additional storage setup: 120 minutes

### Manual Installation Minimum Deployment Time Test Setup

Refer to setup in [OpenStack Installation in Multiple Deployment Models, page -19](#). A multimode deployment model will be used in this test.

### Manual Installation Minimum Deployment Time Procedure

Follow the manual install instructions for COE located here:

[http://docwiki.cisco.com/wiki/Cisco\\_OpenStack\\_Edition:\\_Folsom\\_Manual\\_Install](http://docwiki.cisco.com/wiki/Cisco_OpenStack_Edition:_Folsom_Manual_Install)

Results in installation times of the following (approximately):

- Controller Node: 110 minutes
- Network Node: 55 minutes
- Compute Node: 55 minutes
- Openstack testing, including creation of virtual networks and instances: 40 minutes
- Total time: 260 minutes

### Pass/Fail Criteria

The test fails if the times specified in the test description are not achievable.

## OpenStack Application Infrastructure Deployment Minimum Deployment Time

OpenStack application infrastructure deployment minimum deployment time use case is an application deployment by a system user controlling application deployment times.

An application infrastructure should deploy rapidly. This case is not intended to include application code deployment as part of the definition.

- Deployment of a VM instance from cli or UI: 10 minutes
- Deployment of network from cli or UI: 10 minutes
- Deployment of additional volume from cli or UI: 10 minutes
- Notification of completion of an operation: 10 minutes

### Infrastructure Deployment Minimum Deployment Time Test Setup

A standard multi-node installation of Cisco OpenStack Edition as described in [OpenStack Installation in Multiple Deployment Models, page -19](#) is used for this test.

#### Requirements

- 1 COE Build Node Instance. Dedicated hardware.
- 1 COE Control Node Instance. Dedicated Hardware.
- (at least) 1 COE Compute Node Instance. Dedicated Hardware.
- 1 Cirros testing machine image shall be loaded and available via either installation through Glance and or the horizon dashboard.
- Environment should be pre-qualified for proper operation of python command line client tools and Horizon dashboard.
- No instances shall be running at the start of the test.

### Infrastructure Deployment Minimum Deployment Time Procedure

Perform the following categorical procedures.

#### Deployment from CLI in less than 10 minutes

CLI Deployment scripts are available with elapsed timer for job submission. This alone, however, will not show accurate elapsed time as OpenStack is an asynchronous system. You must watch for proper transition of the instance availability via the nova list command.

- Create instance using the nova boot command and watch status using the nova list command

Follow this simple example for specification of the test procedure:

---

**Step 1** nova boot --image <image name> -- flavor<flavor name> <instance name>.

**Step 2** Check command submission elapsed time and watch for instance to move to an active state by running the nova list.

**Step 3** After a period of 10 minutes, the instance should be in active state.




---

**Note** If an error is encountered or the instance never reaches state 'ACTIVE' then declare test failed.

---

**Deployment of network in less than 10 minutes**

CLI network deployment scripts are available with elapsed timer for job submission. This alone, however, will not show accurate elapsed time as OpenStack is an asynchronous system. You must watch for proper network availability via the quantum net-show command.

- Create network using the quantum net-create command and watch status using the quantum net-show command

Follow this simple example for specification of the test procedure:

- 
- Step 1** quantum net-create private.
- Step 2** Command should complete within 10 minutes.
- Step 3** quantum net-show should list network private.




---

**Note** If an error is encountered running the quantum net-create command, the test result is a failure

---

**Deployment of volume less than 10 minutes**

CLI volume scripts are available with elapsed timer for job submission. This alone, however, will not show accurate elapsed time as OpenStack is an asynchronous system. You must watch for proper volume availability via the cinder list command.

- Create 1GB cinder volume and watch status using the cinder list command

Follow this simple example for specification of the test procedure:

- 
- Step 1** cinder create --display\_name test 1.
- Step 2** Command should complete within 10 minutes.
- Step 3** cinder list.




---

**Note** If an error is encountered running the cinder commands, the test result is a failure.

---

**Pass/Fail Criteria****Instance Create**

Pass: Instance shows ACTIVE via the Nova List Command

Fail: Instance shows status other than ACTIVE, command error or exceeds 10 minutes.

**Network Create**

Pass: quantum net-show command shows the created network

Fail: quantum net-show lists nothing, command error or exceeds 10 minutes.

**Volume Create**

Pass: cinder list command shows the created 1 GB volume

Fail: cinder list command lists nothing, command error or exceeds 10 minutes.