Installation runbook for Huawei Technologies Co., Ltd + Huawei Volume Driver

Partner Name:	Huawei Technologies Co., Ltd
Product Name:	Huawei Volume Driver
Product Version:	Kilo
MOS Version:	7.0
OpenStack version:	Kilo
Product Type:	Cinder Driver

Contents

Document History

Version	Revision Date	Description
0.1	26-1-2016	Initial Version
0.2	3-2-2016	Add switch configuration, hardware list and description
0.3	15-2-2016	Modify the network
0.4	16-2-2016	Add description for switch port
0.5	3-20-2017	Modify the part of cinder configuration
0.3	07-24-2017	Add support for Huawei OceanStor Dorado V3 series storage.

1. Introduction

This document is to serve as a detailed Deployment Guide for Huawei Volume driver with Mirantis OpenStack. Huawei Technologies Co., Ltd offers Huawei Volume Driver to support iSCSI and Fibre Channel connections and enables Huawei Storages such as, OceanStor Dorado V3 series, OceanStor V3 series storage, OceanStor TV2 storage, and OceanStor 18000 high-end storage to provide block storage services for OpenStack cloud environment. This document describes the reference architecture, installation steps for certified MOS + Huawei Volume Driver, limitations and testing procedures.

1.1 Objective

The objective of Mirantis OpenStack certification is to provide Mirantis program partners with an consistent and unified approach for acceptance of their solution into the Mirantis Technology Partner Program.
Certification is designed within the context of Mirantis OpenStack infrastructure, including Mirantis Fuel deployment tool and supported cloud reference architectures.

1.2 Target Audience

This documentation is intended for OpenStack administrators provided with the details of the setup, configuration and tests run on Mirantis OpenStack with Huawei Volume driver for the purpose of Unlocked OpenStack Driver Validation. It assumes that you have experience with network and cloud concepts.

2. Product Overview

Huawei Volume Driver enables integration between OpenStack Cloud and Huawei Storages such as, OceanStor Dorado V3 series, OceanStor V3 series storage, OceanStor TV2 storage, and OceanStor 18000 high-end storage to provide block storage services for OpenStack.
Huawei Volume Driver implemented OpenStack Cinder accesses Huawei storage service APIs, such as volume, snapshot, Qos etc and supports various operations.

3. Joint reference architecture

Huawei Volume Driver architecture. Prints in red modules are implemented by Huawei Volume Driver.

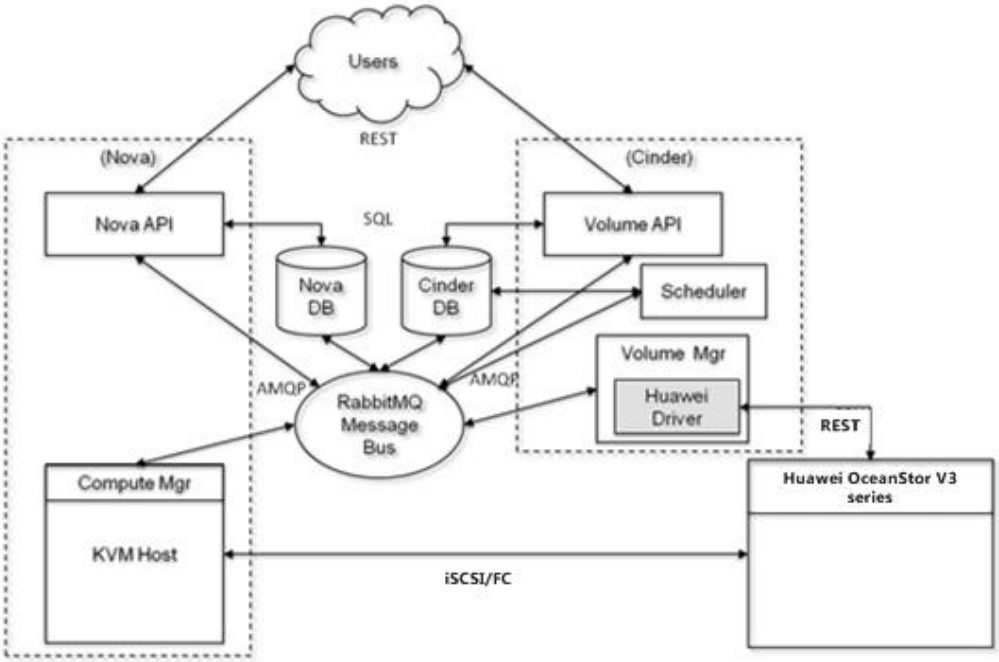

4. Networking

4.1 Network Topology

It is highly recommended that the storage network is broken out as either a physical interface, or as a separate VLAN, to keep the traffic segregated from other traffic types. If you have enough physical interfaces, we strongly recommend using a separate physical interface for storage network.

Each Server at least 2 NICs are required.

- eth0 - untagged port for PXE network (Fuel network), Management (VLAN 101), Public/Floating (VLAN 102)
- eth1 - port for networks: Storage (untagged)

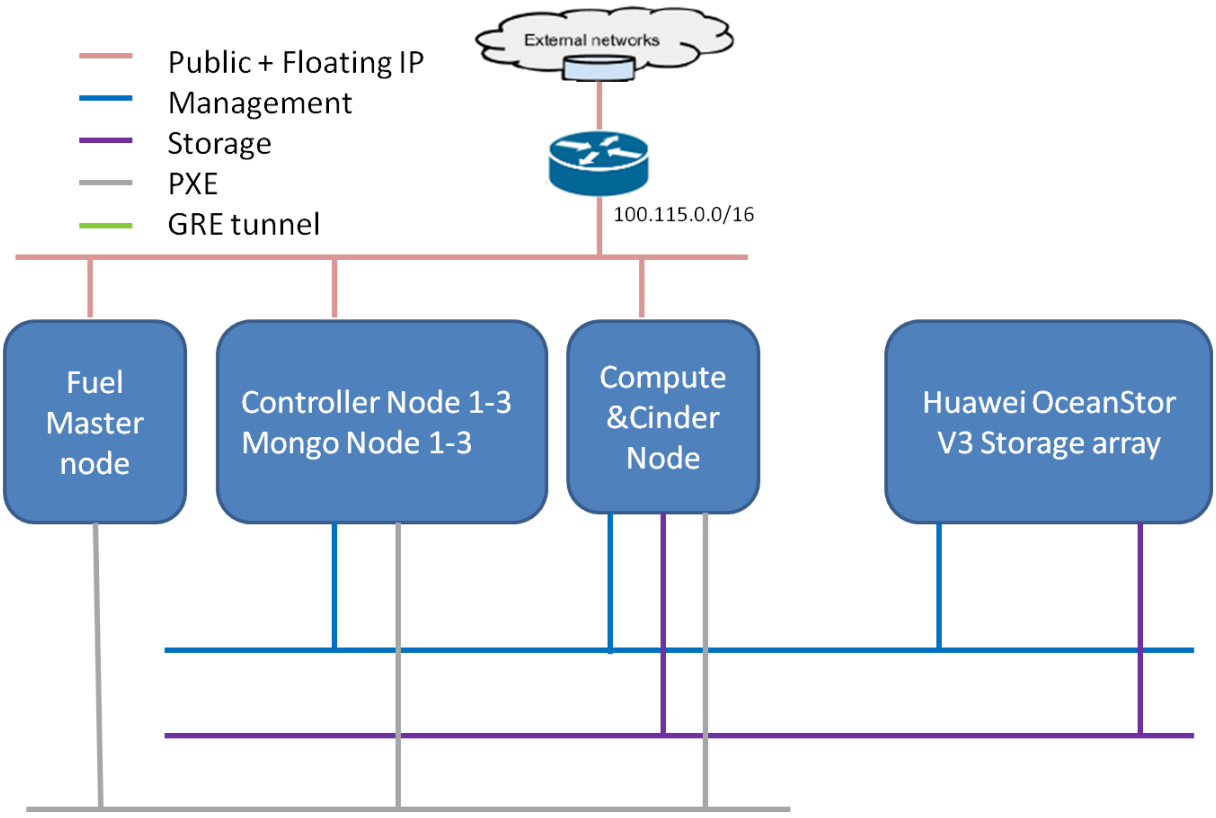

5. Installation and Configuration

5.1 Overview of MOS installation steps

- ◇ Setup OpenStack Network configuration in Switch device.
- ◇ Install Fuel Master Node on Physical Server. Customer can use a VM for fuel master by following the MOS user guide. In this demo, we use a physical server instead.
- ● Deploy MOS Slave nodes on Fuel UI
- ◇ Run Fuel Health Check, make sure all cases pass.
- ◇ Manually integrated Huawei OceanStor V3 storage and Huawei Volume driver with MOS.
- ◇ Run Fuel Heath Check test on Fuel UI
- ◇ Verify Huawei Volume Driver basic functions on MOS, plus additional manual/automated acceptance tests.

5.2 MOS Installation in details

5.2.1 Hardware and environment

The following hardware devices are used in this demo.

Name	Device	Count
Server	Huawei RH2285-v2	3
	Huawei RH2285	2
Network Switch	Huawei Quidway S3300	1
Storage	Huawei OceanStor 5800 V3	1

5.2.2 Switch configuration

Use the following configuration to deploy Mirantis OpenStack using Huawei Quidway S3300 network switch. For one master node and four Openstack nodes of physical servers, each server have two NICs with following VLAN IDs are used:

- eth0 - untagged port for PXE network (Fuel network), Management (VLAN 101), Public/Floating (VLAN 102)
- eth1- port for networks: Storage (untagged)

The connection between switch port and server NIC as below: each switch port used for one sever NIC.

Switch Port	Server name	Server NIC
E0/0/1	Fuel Master	eth0
E0/0/2	Controller Node 1	eth0
E0/0/3	Controller Node 2	eth0
E0/0/4	Controller Node 3	eth0
E0/0/5	Compute&Cinder Node 1	eth0
E0/0/6	Fuel Master	eth1
E0/0/7	Controller Node 1	eth1
E0/0/8	Controller Node 2	eth1
E0/0/9	Controller Node 3	eth1
E0/0/10	Compute&Cinder Node 1	eth1

The details port configuration for Huawei Quidway S3300 network switch is below:

```
#
!Software Version V100R005C01SPC100
sysname openstack_sw
#
vlan batch 101 to 102 1000 to 1030
#
vlan 1
interface Vlanif1
ip address 100.115.114.254 255.255.0.0
#
interface Ethernet0/0/1
port link-type trunk
port trunk allow-pass vlan 101 to 102 1000 to 1030
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/2
port link-type trunk
port trunk allow-pass vlan 101 to 102 1000 to 1030
ntdp enable
ndp enable
bpdu enable
dhcp snooping trusted
#
interface Ethernet0/0/3
port link-type trunk
port trunk allow-pass vlan 101 to 102 1000 to 1030
ntdp enable
ndp enable
bpdu enable
#
```

interface Ethernet0/0/4
port link-type trunk
port trunk allow-pass vlan 101 to 102 1000 to 1030
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/5
port link-type trunk
port trunk allow-pass vlan 101 to 102 1000 to 1030
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/6
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/7
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/8
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/9
ntdp enable
ndp enable
bpdu enable
#
interface Ethernet0/0/10
ntdp enable
ndp enable
bpdu enable

- **For Fuel Master Node**

◇ Mount MirantisOpenStack-7.0.iso through IPMI on one Huawei RH2285-V2 server and power the machine on. Set the boot order for the system with the Huawei DVD media.
Using the Fuel User guide (https://docs.mirantis.com/openstack/fuel/fuel-7.0/user-guide.html#download-and-install-fuel) go through the installation process until Fuel Setup will appear. Enter into it.

◇ Navigate to Network Setup menu, adjust configuration to one provided in the table below and then click 'Check' button to ensure the configuration is valid.

Interface	IP address	Netmask	Default Gateway
eth0	100.115.114.230	255.255. 0.0	100.115.0.1

eth1(PXE)	192.168.10.10	255.255.255.0	-

◇ Navigate to PXE Setup menu, adjust configuration to one provided in the table below and then click 'Check' button to ensure the configuration is valid.

Interface	DHCP Pool Start	DHCP Pool End	DHCP Gateway
eth1(PXE)	192.168.10.110	192.168.10.254	192.168.10.10

◇ Navigate to DNS & Hostname menu, adjust configuration to one provided in the table below and then click 'Check' button to ensure the configuration is valid. It will also check that the Internet is reachable by the Fuel node because it will be required for the MOS installation later.

Hostname	Domain	Search Domain	External DNS	Hostname to test DNS
fuel	Domain.tld	Domain.tld	8.8.4.4, 8.8.8.8	baidu.com

◇ Keep Time Sync menu untouched with default settings entered. Click to 'Check' button to ensure NTP servers are reachable.
◇ Change default root password to make sure Fuel Health Check cases pass.
◇ Save and Quit Setup and boot the Fuel Master Node.
◇ When the installation is done Fuel UI become available at https://100.115.114.230:8443 as below.

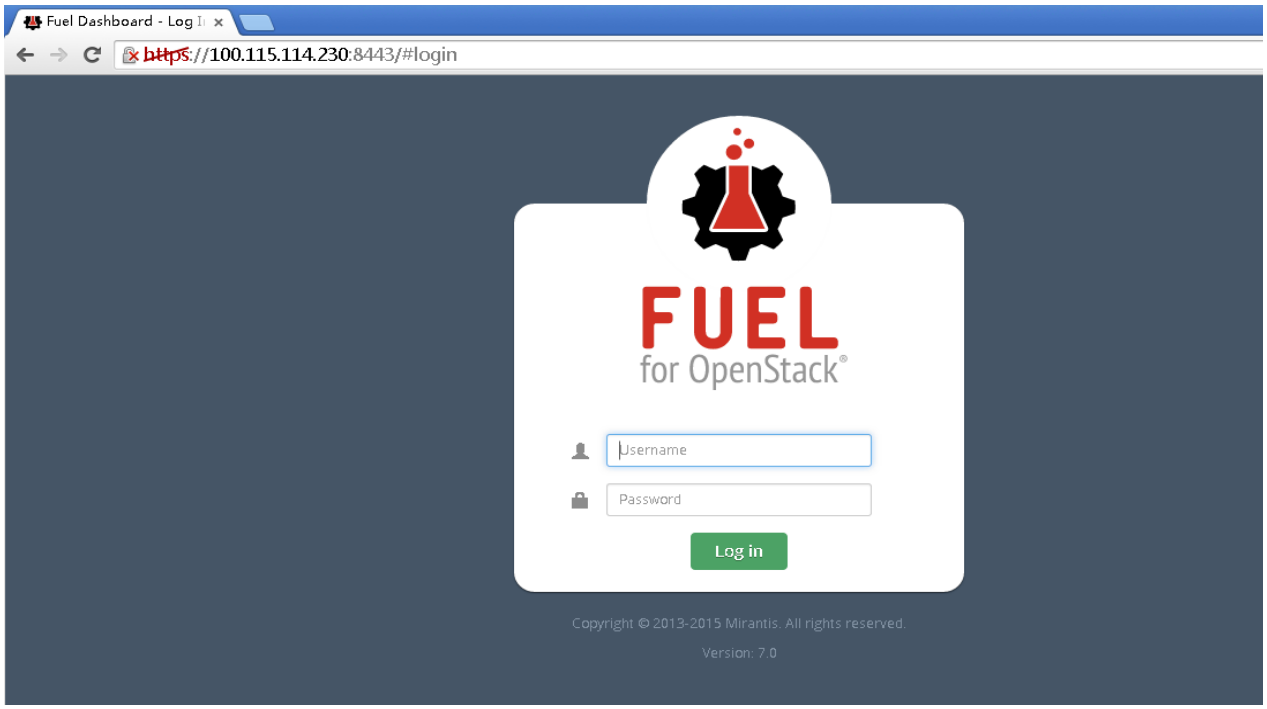

5.3 Creation of OpenStack environment

The settings for the OpenStack cloud in this demo are below:

- Name: Cert
- OpenStack re: Kilo on Ubuntu 14.04
- Deployment mode: HA

- Compute: KVM
- Cinder backend: Cinder LVM over iSCSI for volumes (Use default providers)
- Glance backend: Swift (Use default providers)
- Network setup: Neutron with VLAN segmentation
- Ceilometer: Install
- Murano and Sahara: Not Install, please install them if need.

Note:

For creating OpenStack environment, please find the detail description in MOS user guide: https://docs.mirantis.com/openstack/fuel/fuel-7.0/user-guide.html#create-a-new-openstack-environment

1) On the Fuel UI, click on “New OpenStack Environment”.
2) When the wizard opens, enter the name and the desired OpenStack Release.

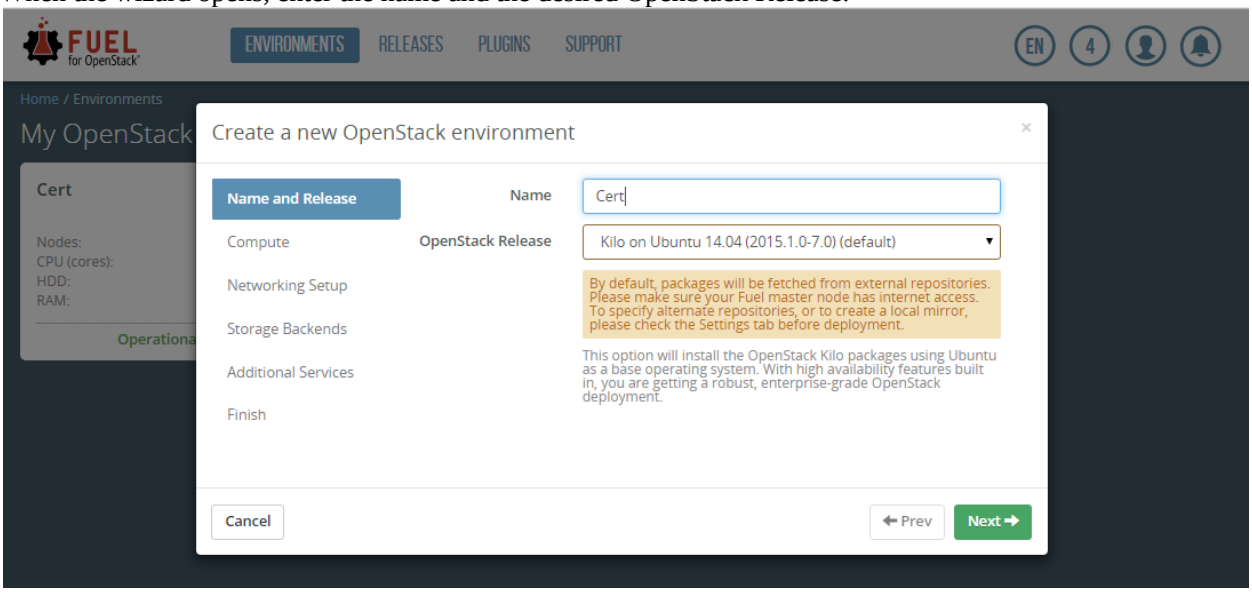

3) Select the Compute for the Environment

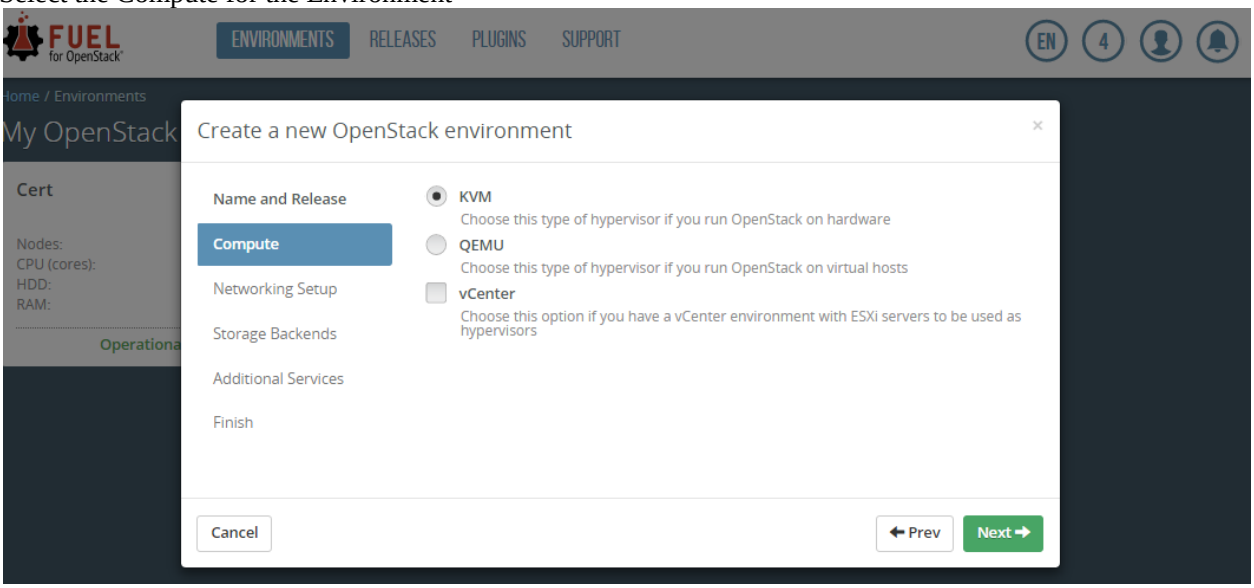

4) Select the required Neutron Setup.

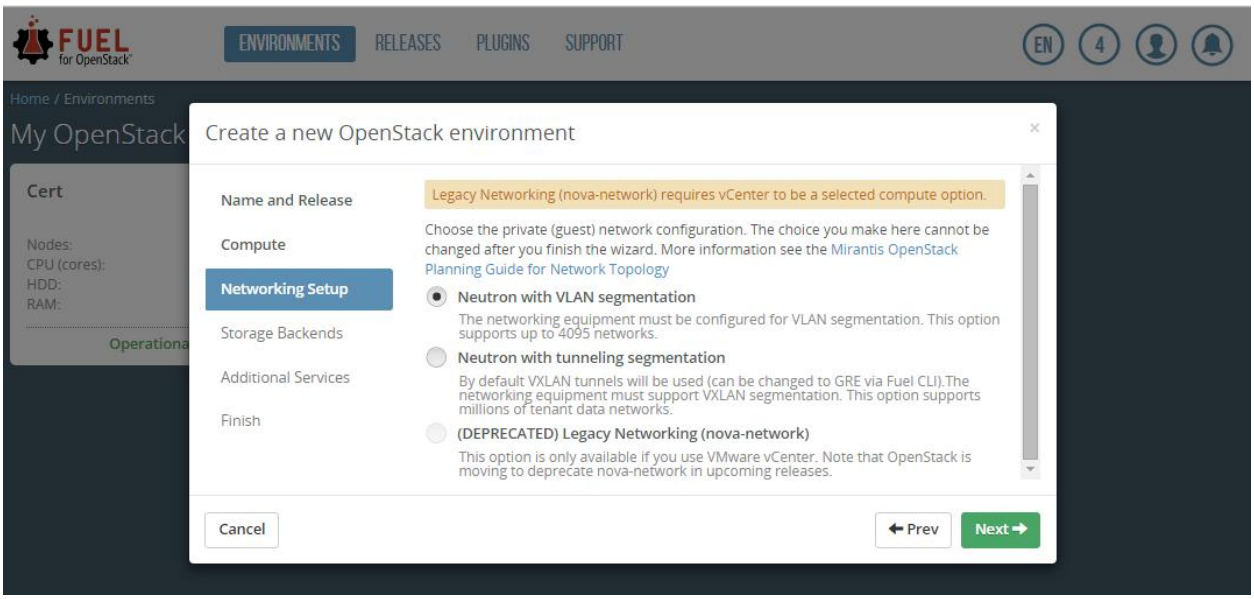

5) Under Storage Backends, leave the default option. Huawei Volume driver can be installed after the OpenStack is deployed.

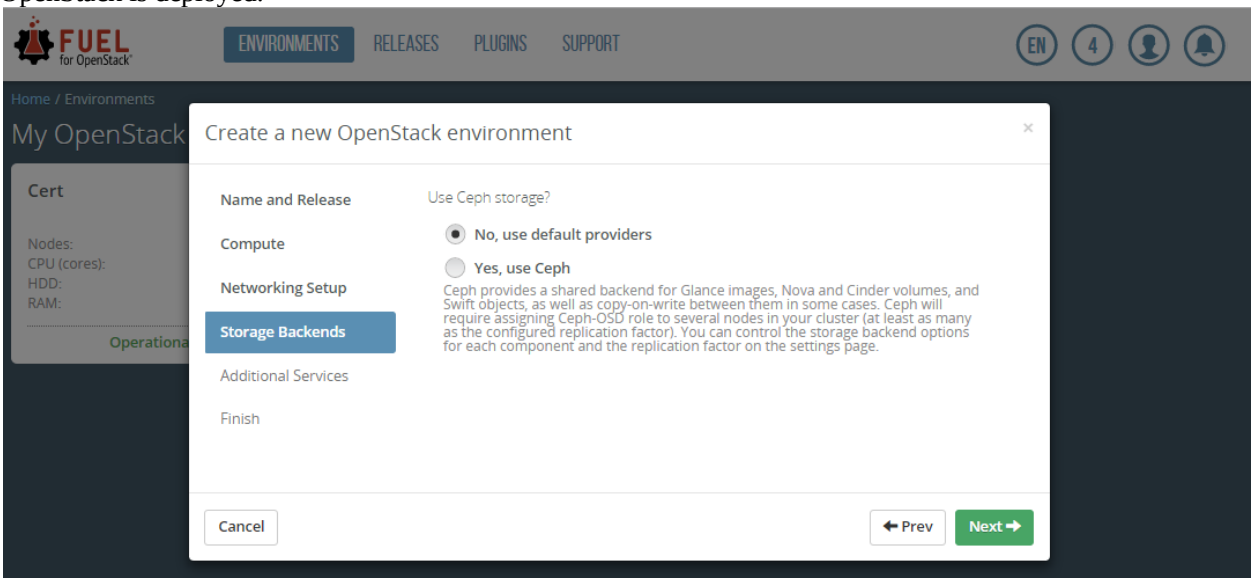

6) Select the additional services and click on next.

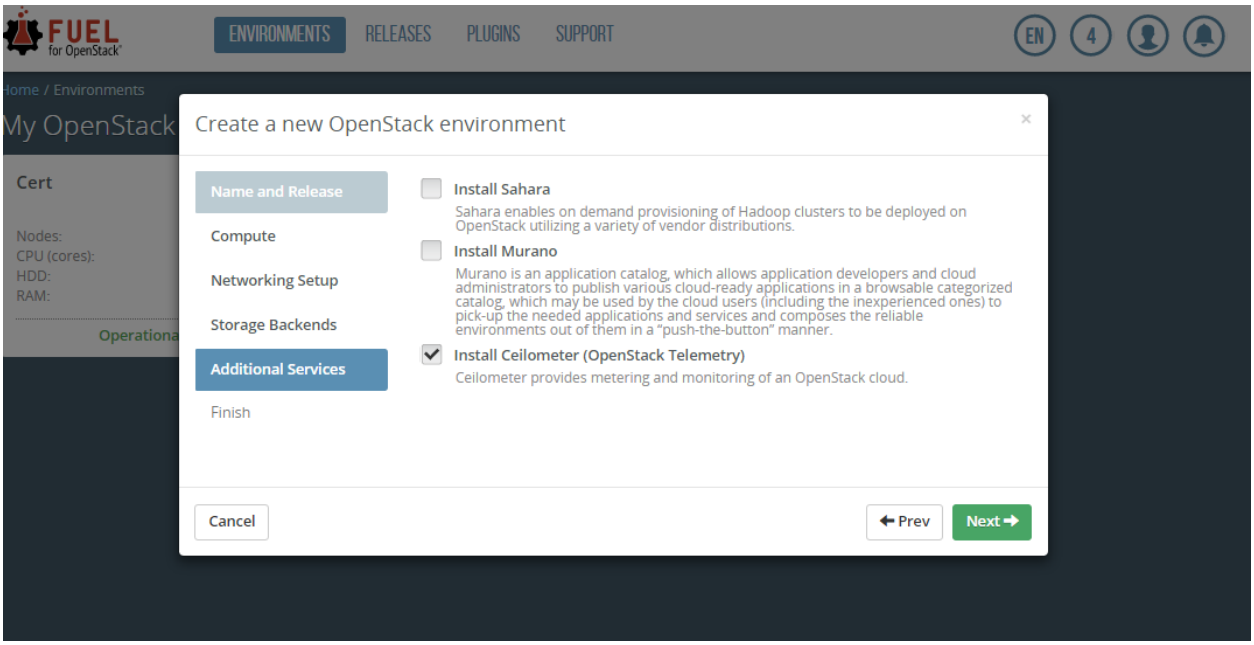

7) Click Create to start deploy the OpenStack.

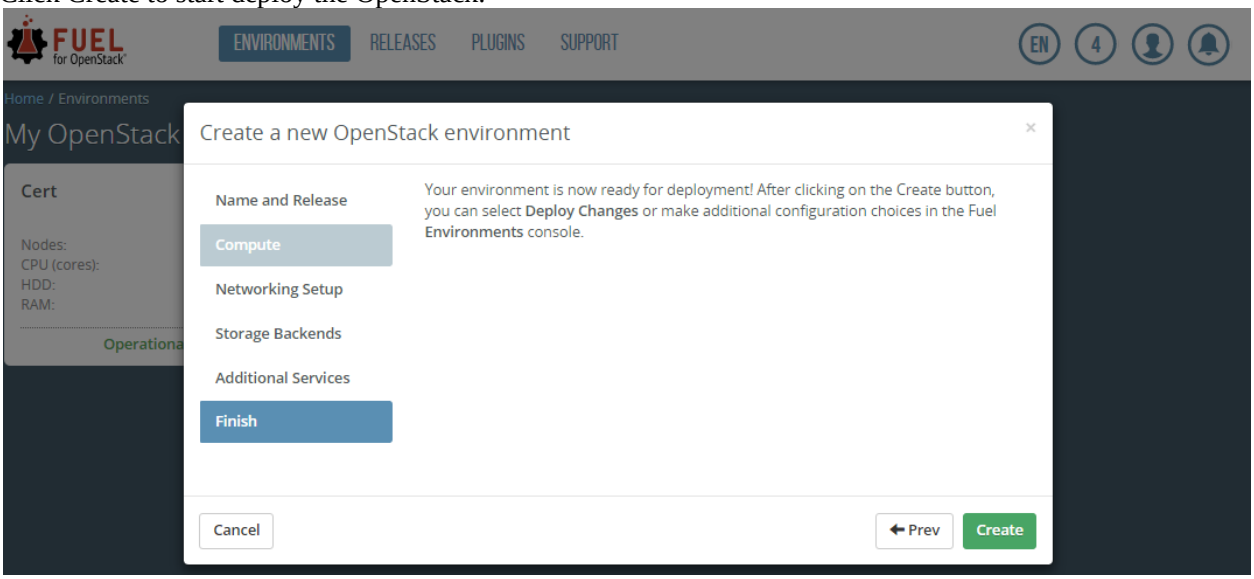

5.4 MOS Deployment

5.4.1 Nodes

All nodes: 3 Controller nodes combined with Telemetry-Mongo DB, and 1 Compute node combined with Storage-Cinder.

For this demo, hardware model information as follow:

Controller Node 1	Controller Node 2	Controller Node 3	Compute Node
RH2285-v2	RH2285	RH2285	RH2285-v2

Once all the nodes have booted up through PXE, they appear on the Fuel UI. Configure controllers, cinder and compute.

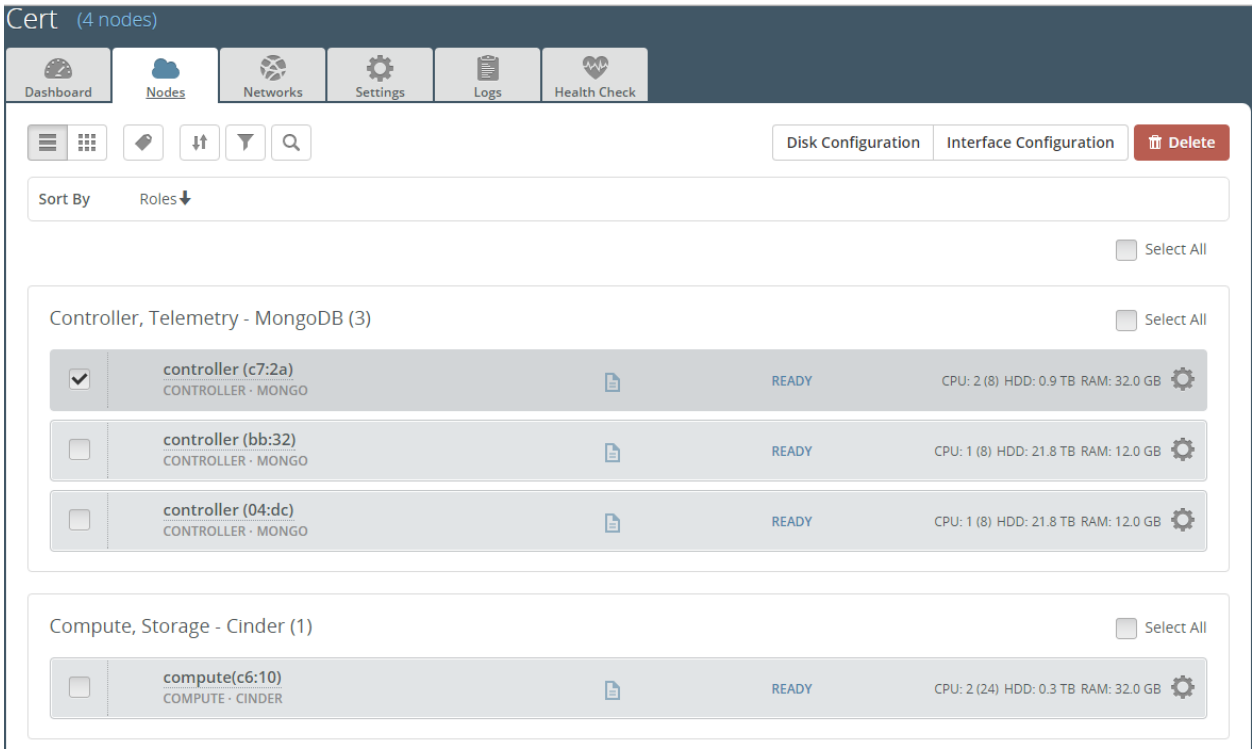

Controller node (c7:2a) disks configuration:

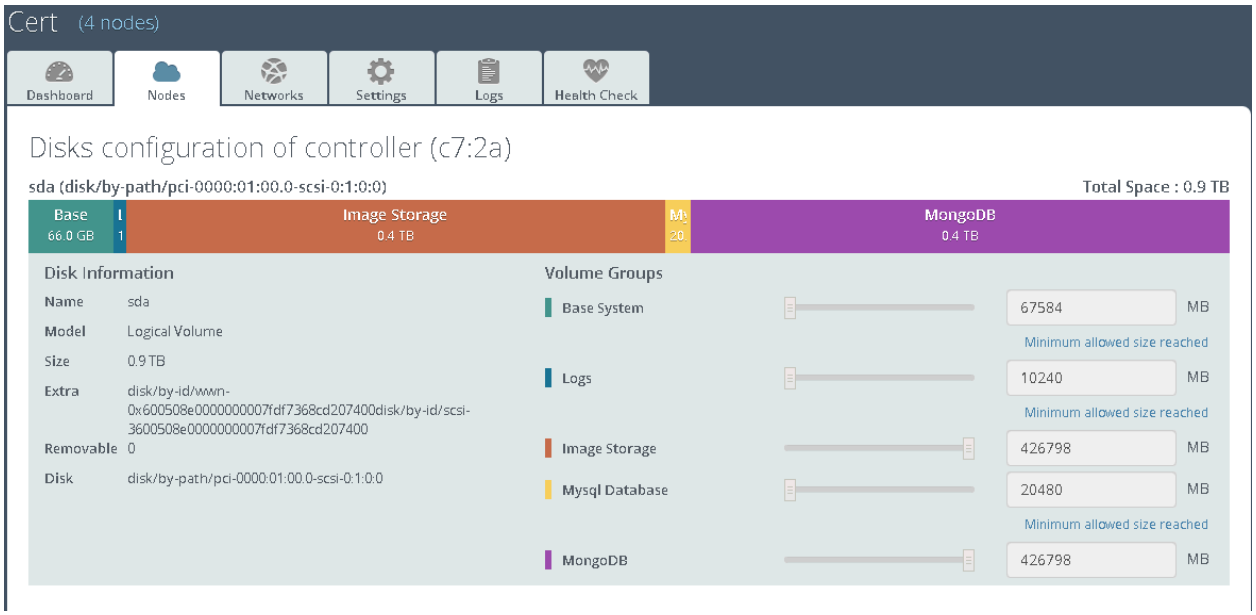

Controller node (c7:2a) network Interfaces configuration:

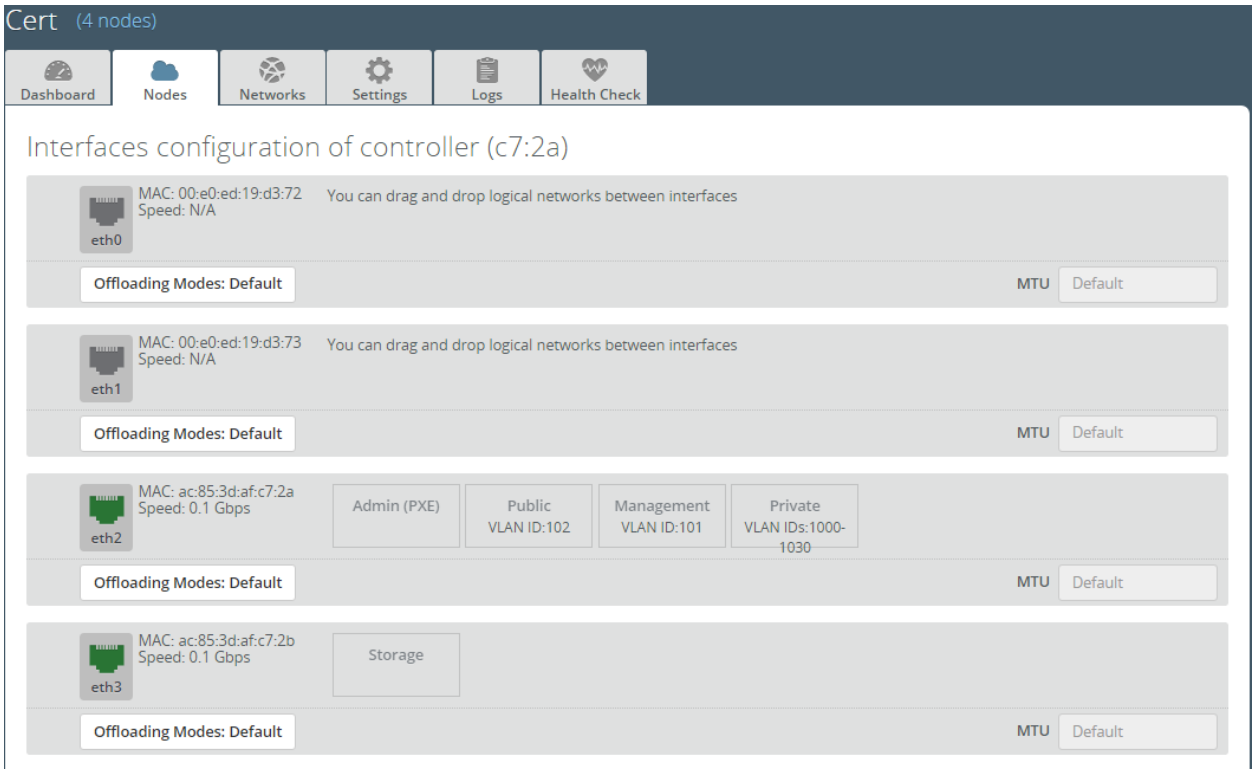

Compute Node (c6:10) disks configuration:

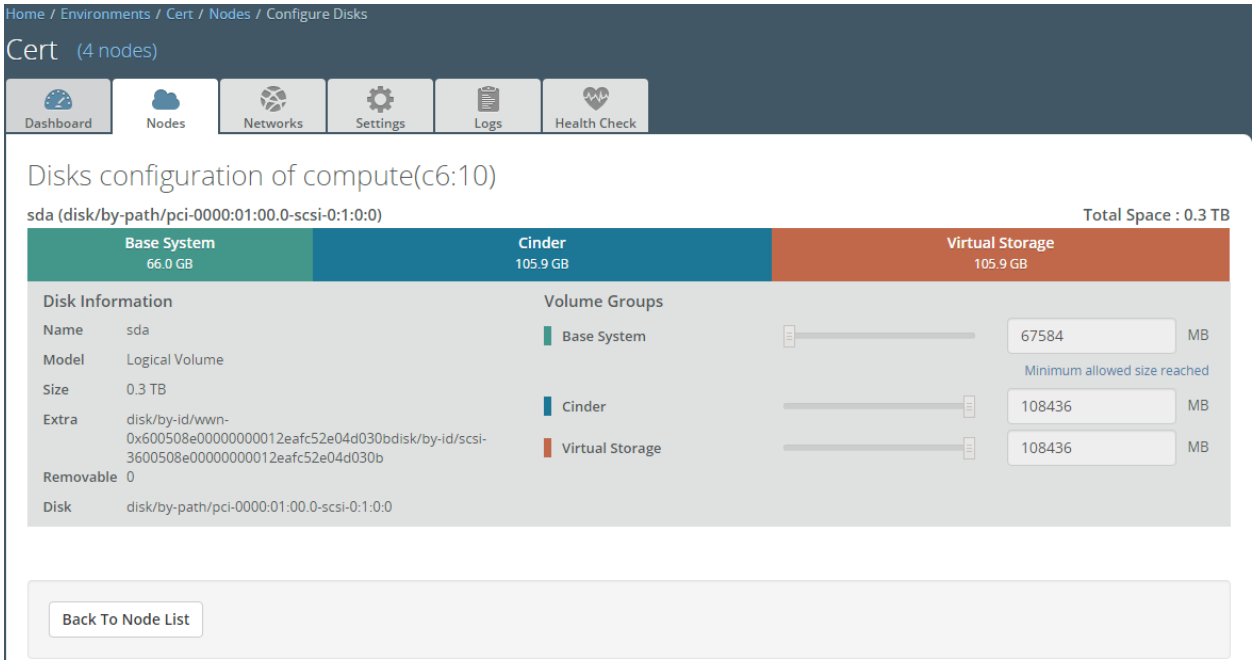

Compute Node (c6:10) network interfaces configuration:

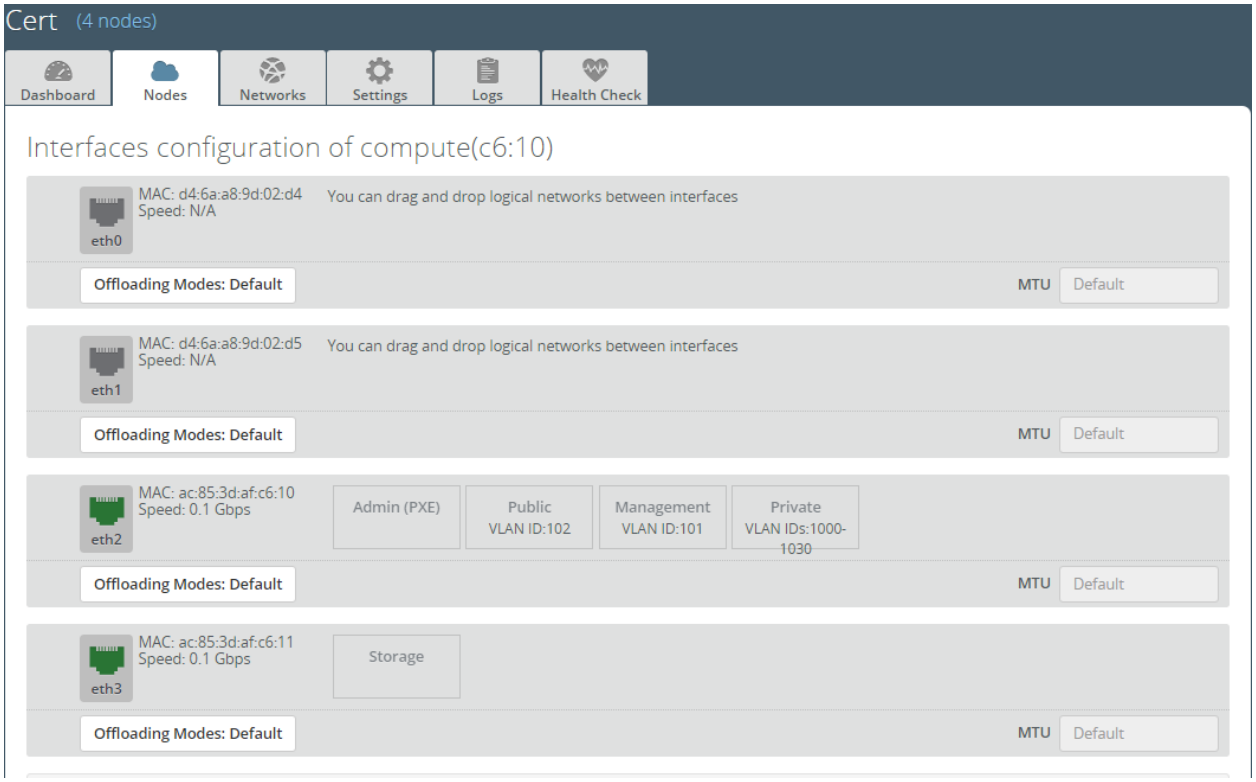

5.4.2 Network Setting

Change the Public, Storage, Management and Private network as your plan. After configuring the nodes, run "Verify Networks" under Networks tab. The test should pass.

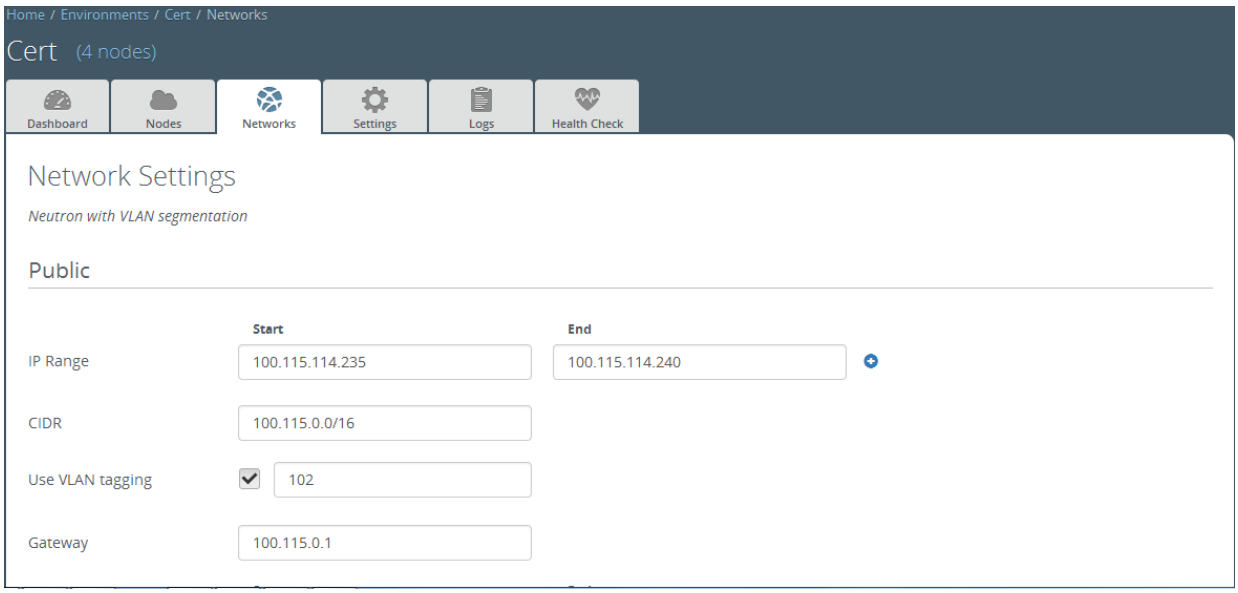

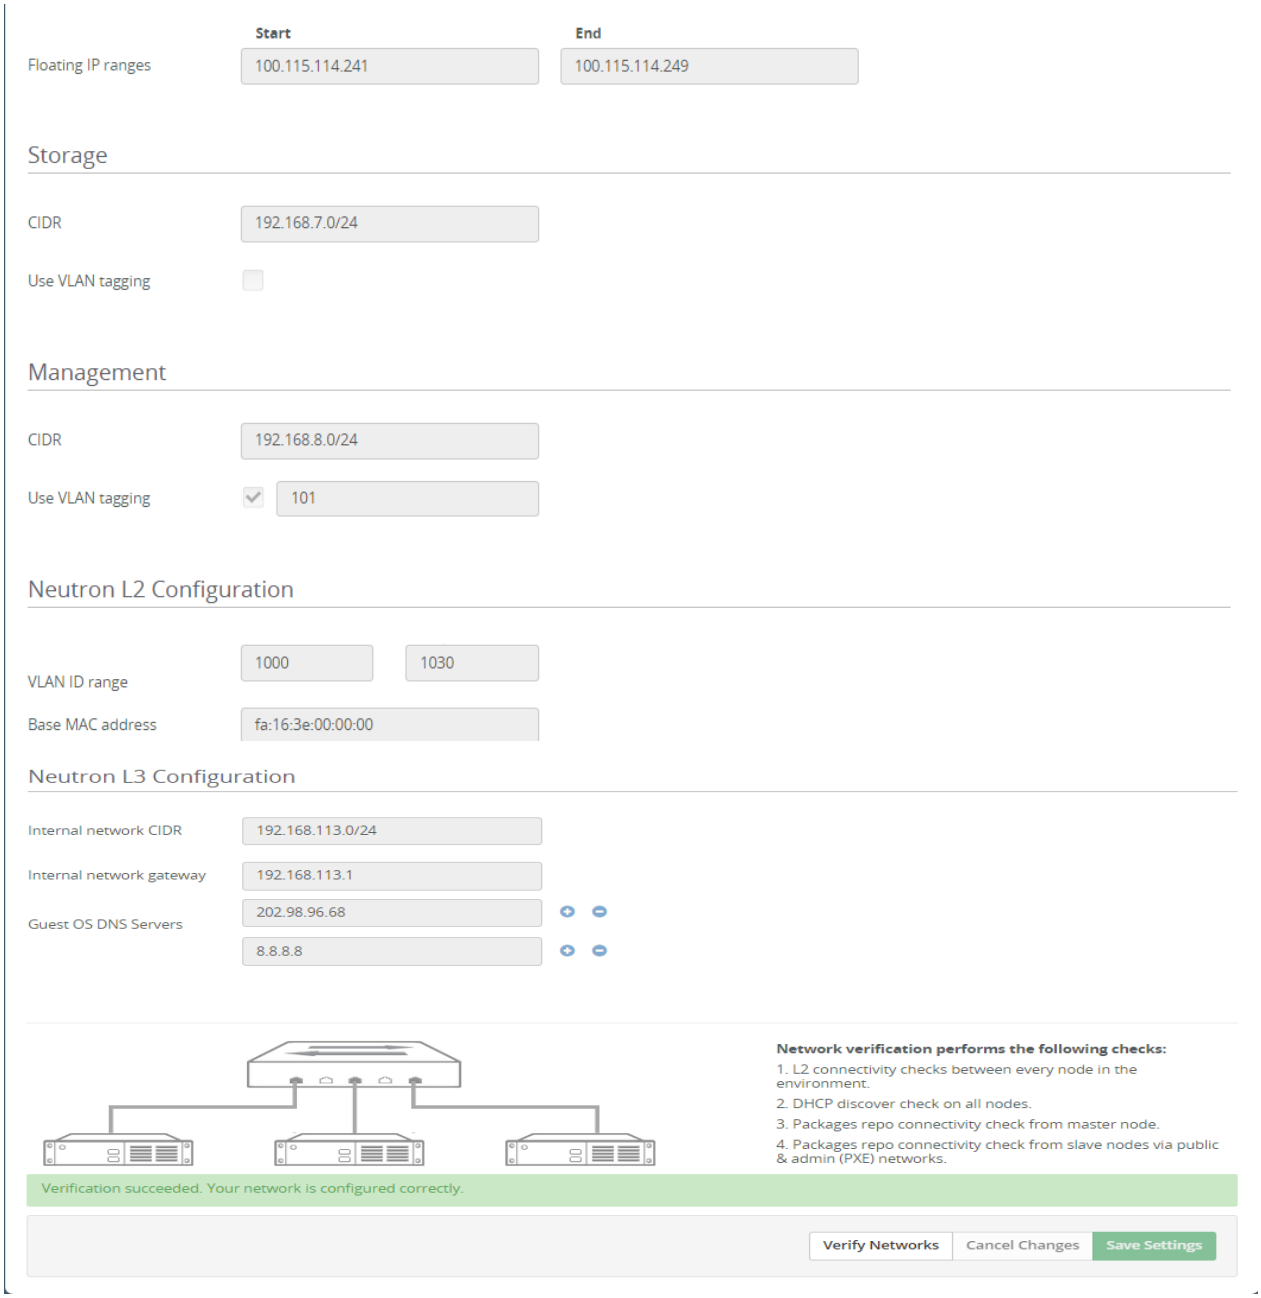

5.4.3 OpenStack Settings

The "Settings" tab allows you to set or modify various values for the system. Many other values can be set by editing configuration files and running command-line tools on the nodes. You can find the details description in user guide:
https://docs.mirantis.com/openstack/fuel/fuel-7.0/user-guide.html#settings-tab

In this demo, we change the default horizon setting, DNS and NTP setting as blew. The other settings are keep default.

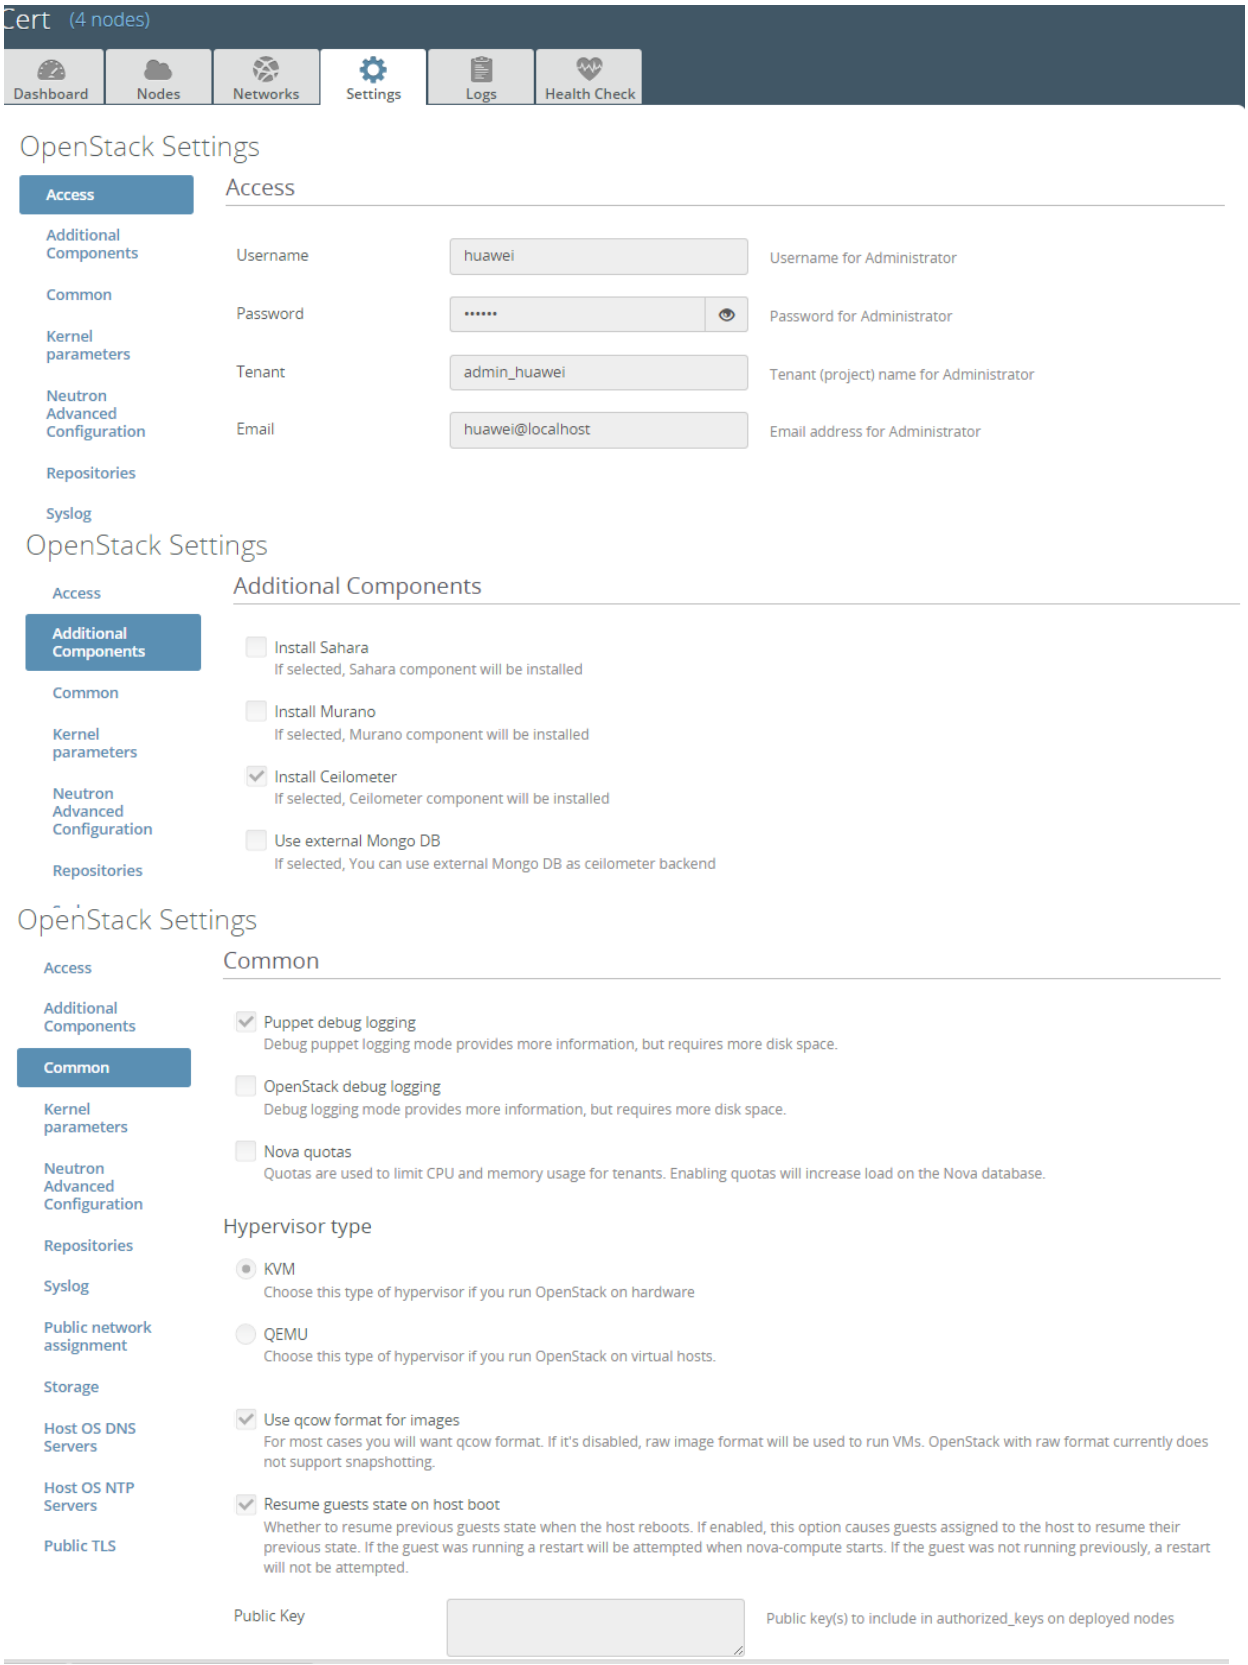
Cert (4 nodes)
Dashboard
Nodes
Networks
Settings
Logs
Health Check
OpenStack Settings
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Access
Username
huawei
Username for Administrator
Password
Password for Administrator
Tenant
admin_huawei
Tenant (project) name for Administrator
Email
huawei@localhost
Email address for Administrator
OpenStack Settings
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Additional Components
Install Sahara
If selected, Sahara component will be installed
Install Murano
If selected, Murano component will be installed
Install Ceilometer
If selected, Ceilometer component will be installed
Use external Mongo DB
If selected, You can use external Mongo DB as ceilometer backend
OpenStack Settings
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Public network assignment
Storage
Host OS DNS Servers
Host OS NTP Servers
Public TLS
Common
Puppet debug logging
Debug puppet logging mode provides more information, but requires more disk space.
OpenStack debug logging
Debug logging mode provides more information, but requires more disk space.
Nova quotas
Quotas are used to limit CPU and memory usage for tenants. Enabling quotas will increase load on the Nova database.
Hypervisor type
KVM
Choose this type of hypervisor if you run OpenStack on hardware
QEMU
Choose this type of hypervisor if you run OpenStack on virtual hosts.
Use qcow format for images
For most cases you will want qcow format. If it's disabled, raw image format will be used to run VMs. OpenStack with raw format currently does not support snapshotting.
Resume guests state on host boot
Whether to resume previous guests state when the host reboots. If enabled, this option causes guests assigned to the host to resume their previous state. If the guest was running a restart will be attempted when nova-compute starts. If the guest was not running previously, a restart will not be attempted.
Public Key
Public key(s) to include in authorized_keys on deployed nodes

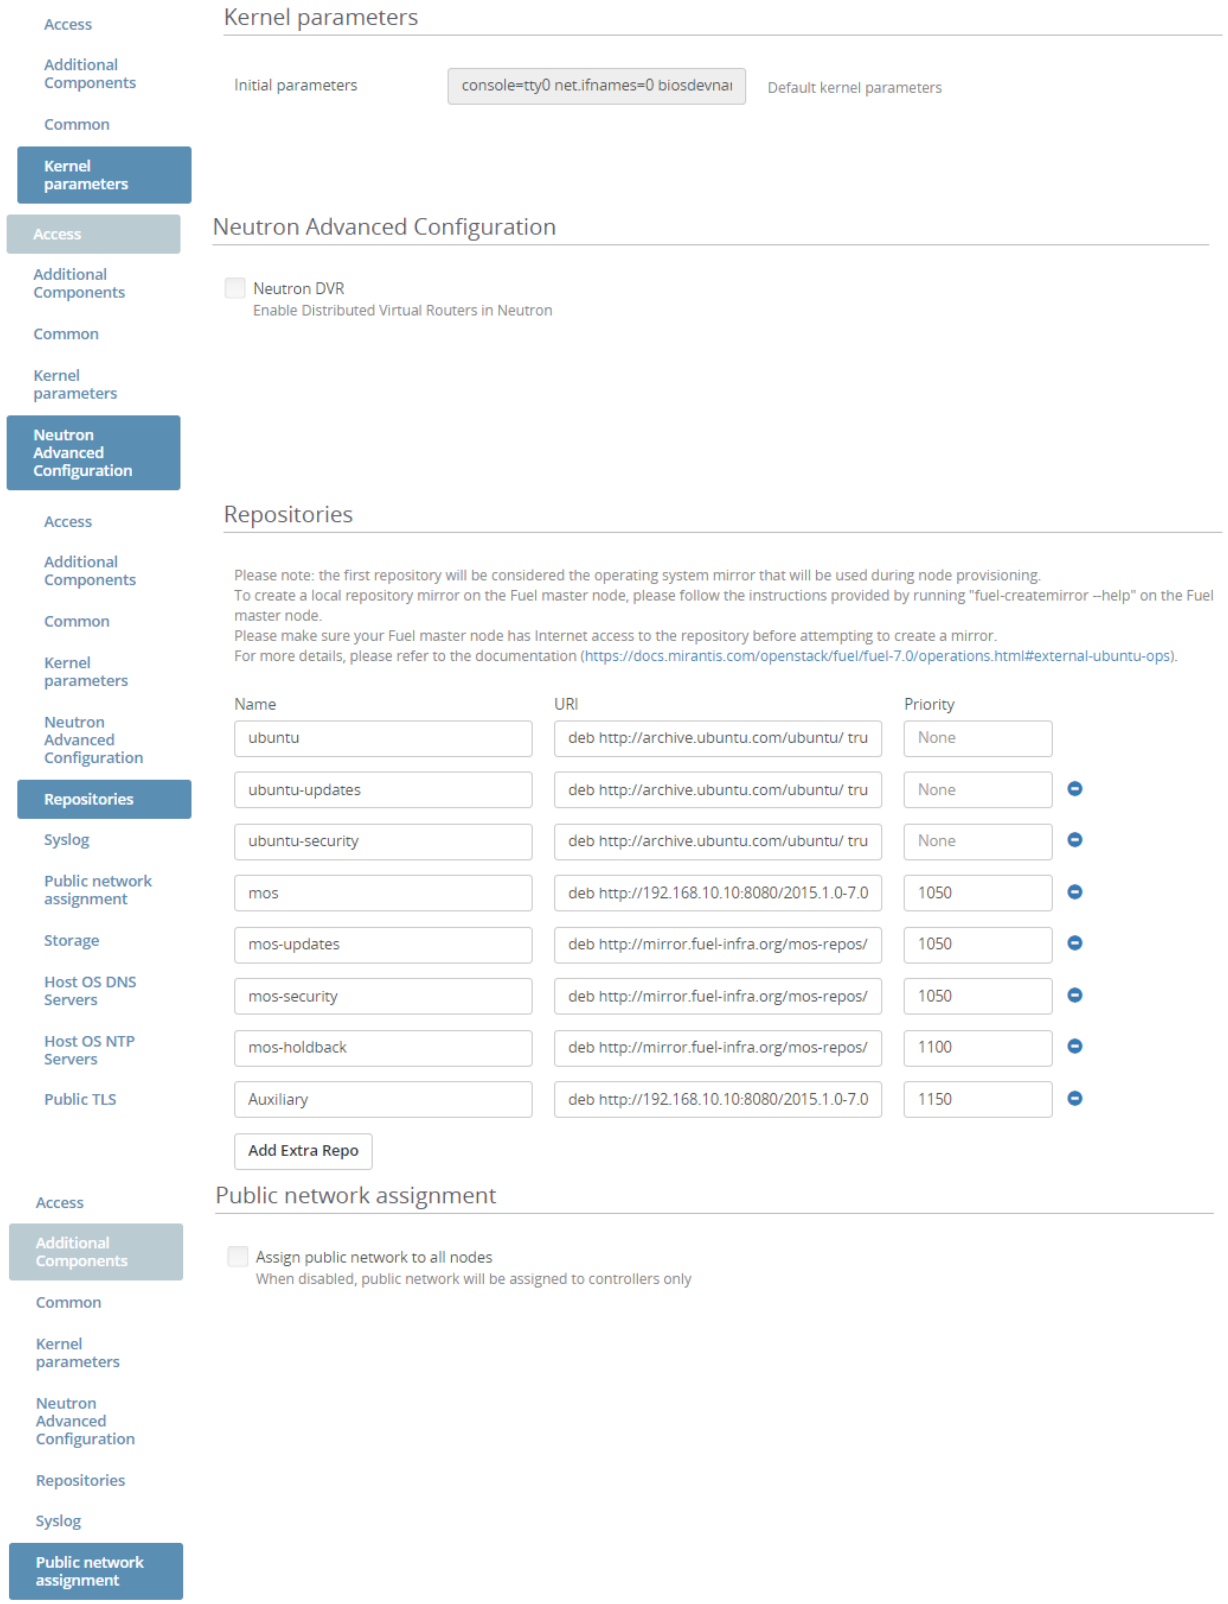
Access
Additional Components
Common
Kernel parameters
Kernel parameters
Initial parameters
console=tty0 net.ifnames=0 biosdevnai
Default kernel parameters
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Neutron Advanced Configuration
Neutron DVR
Enable Distributed Virtual Routers in Neutron
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Public network assignment
Storage
Host OS DNS Servers
Host OS NTP Servers
Public TLS
Repositories
Please note: the first repository will be considered the operating system mirror that will be used during node provisioning.
To create a local repository mirror on the Fuel master node, please follow the instructions provided by running "fuel-createmirror --help" on the Fuel master node.
Please make sure your Fuel master node has Internet access to the repository before attempting to create a mirror.
For more details, please refer to the documentation (https://docs.mirantis.com/openstack/fuel/fuel-7.0/operations.html#external-ubuntu-ops).
Name
URI
Priority
ubuntu
deb http://archive.ubuntu.com/ubuntu/ tru
None
ubuntu-updates
deb http://archive.ubuntu.com/ubuntu/ tru
None
ubuntu-security
deb http://archive.ubuntu.com/ubuntu/ tru
None
mos
deb http://192.168.10.10:8080/2015.1.0-7.0
1050
mos-updates
deb http://mirror.fuel-infra.org/mos-repos/
1050
mos-security
deb http://mirror.fuel-infra.org/mos-repos/
1050
mos-holdback
deb http://mirror.fuel-infra.org/mos-repos/
1100
Auxiliary
deb http://192.168.10.10:8080/2015.1.0-7.0
1150
Add Extra Repo
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Public network assignment
Public network assignment
Assign public network to all nodes
When disabled, public network will be assigned to controllers only

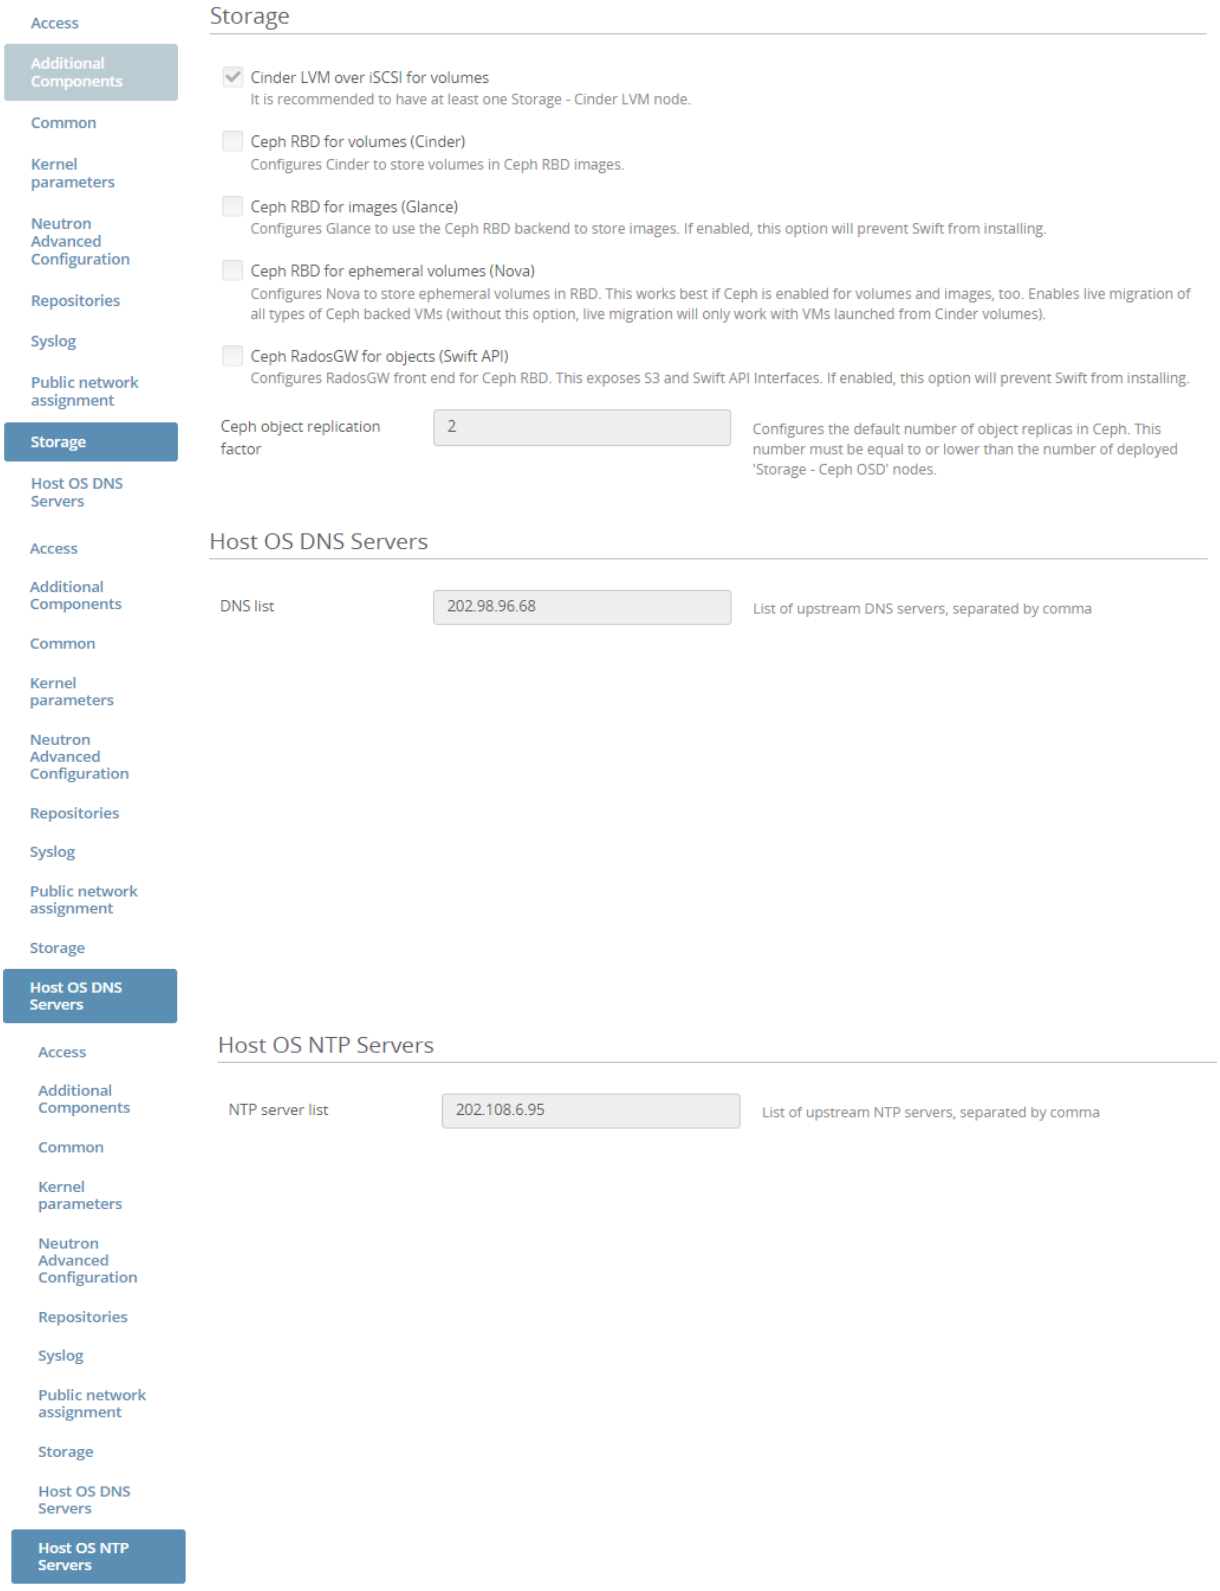
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Public network assignment
Storage
Host OS DNS Servers
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Public network assignment
Storage
Host OS DNS Servers
Access
Additional Components
Common
Kernel parameters
Neutron Advanced Configuration
Repositories
Syslog
Public network assignment
Storage
Host OS DNS Servers
Host OS NTP Servers
Storage
Cinder LVM over iSCSI for volumes
It is recommended to have at least one Storage - Cinder LVM node.
Ceph RBD for volumes (Cinder)
Configures Cinder to store volumes in Ceph RBD images.
Ceph RBD for images (Glance)
Configures Glance to use the Ceph RBD backend to store images. If enabled, this option will prevent Swift from installing.
Ceph RBD for ephemeral volumes (Nova)
Configures Nova to store ephemeral volumes in RBD. This works best if Ceph is enabled for volumes and images, too. Enables live migration of all types of Ceph backed VMs (without this option, live migration will only work with VMs launched from Cinder volumes).
Ceph RadosGW for objects (Swift API)
Configures RadosGW front end for Ceph RBD. This exposes S3 and Swift API Interfaces. If enabled, this option will prevent Swift from installing.
Ceph object replication factor
2
Configures the default number of object replicas in Ceph. This number must be equal to or lower than the number of deployed 'Storage - Ceph OSD' nodes.
Host OS DNS Servers
DNS list
202.98.96.68
List of upstream DNS servers, separated by comma
Host OS NTP Servers
NTP server list
202.108.6.95
List of upstream NTP servers, separated by comma

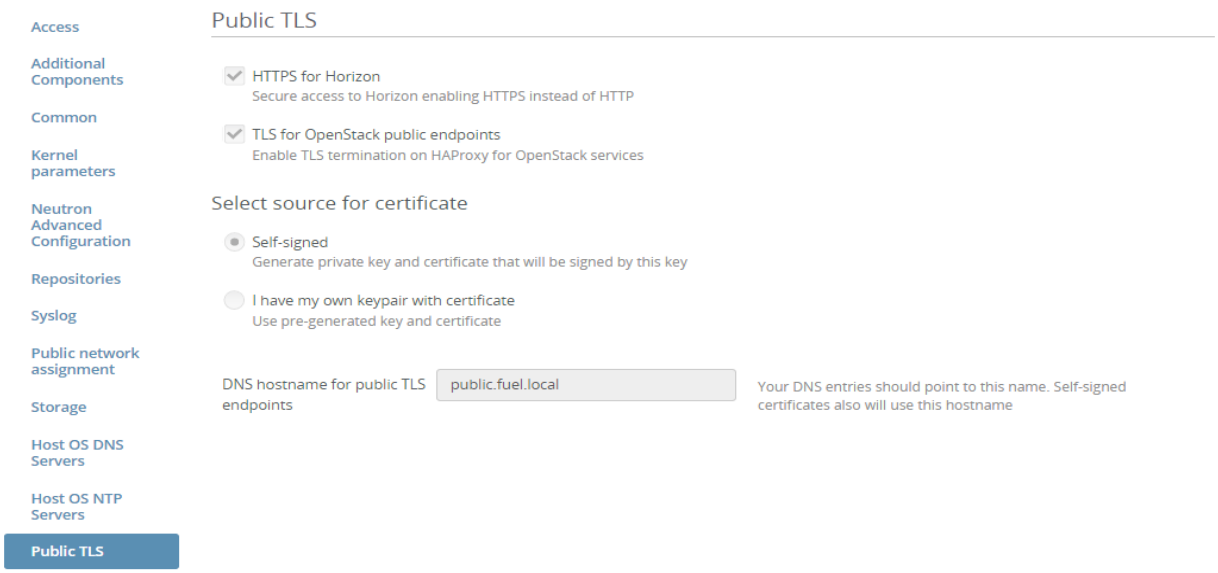

5.4.4 The followings are bridge names and their IP addresses after the deployment for the reference:

Network Interface	Controller (combine Mongo)Node 1	Controller (combine Mongo)Node 2	Controller (combine Mongo)Node 3	Compute (combine Cinder) Node
eth0 (br-fw-admin)	192.168.10.110	192.168.10.111	192.168.10.117	192.168.10.116
eth1(br-ex/public)	IP: 100.115.114.238 VIP:100.115.114.236	100.115.114.237	100.115.114.239	-
eth1 (br-mgnt)	192.168.8.5	192.168.8.3	192.168.8.4	192.168.8.6
eth1 (br-storage)	192.168.7.4	192.168.7.2	192.168.7.3	192.168.7.1

5.5 Huawei Volume Driver installation steps

The following is the procedure that can be followed to install the Huawei Volume Driver. Make sure all cable connection between Compute node and Huawei OceanStor V3 storage are finished. In this demo, we use fibre to connect computer node and storage device directly.

Note:

In production environment, Huawei still recommends to using a dedicated SAN network for OpenStack environment.

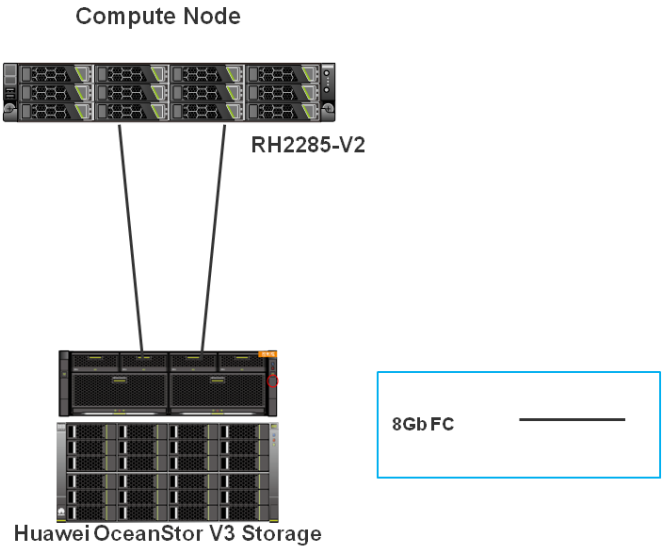

1) Hardware Preparation:
   - Huawei OceanStor V3 series storage (Huawei OceanStor 5800 V3 is used in this demo)

2) Software Preparation:
   - Huawei Volume Driver

5.5.1 Get Huawei Volume Driver

You can get the Huawei Volume Driver for Huawei Storage from GitHub website in below link:
https://github.com/huaweistorage/OpenStack_Driver/tree/master

Unzip OpenStack_Driver-master.rar , it has the following 1 folder

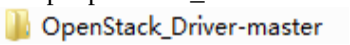

Inside of the folder it has two folders

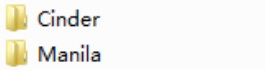

Inside of the Cinder, there are five versions of the OpenStack,

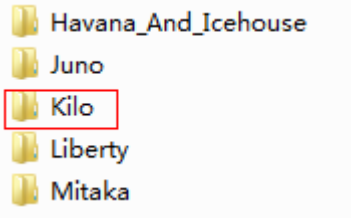

5.5.2 Cinder configuration on all cinder nodes

In this demo, Cinder component is installed in compute node. The following configuration is done on compute node.

1) Please create a folder names '*huawei*' in path */usr/lib/python2.7/dist-packages/cinder/volume/drivers/*. If it exists, backup this folder and remove all files that already in *huawei* folder.

2) using SCP to upload all files from above *kilo* folder to */usr/lib/python2.7/dist-packages/cinder/volume/drivers/huawei/*.

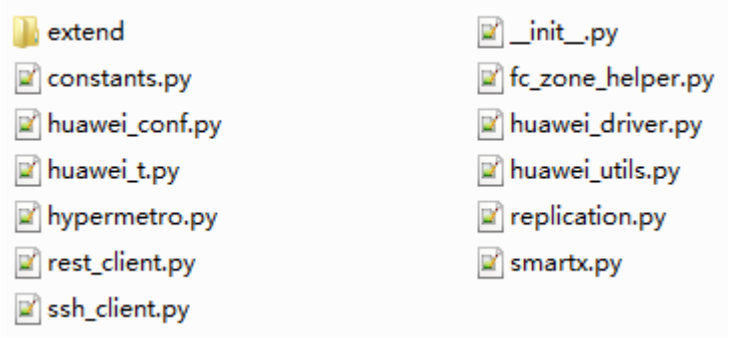

3) Configuring the volume driver for V3 (Fibre Channel)

This configuration file is an example for this certification scenario, for details please checks below description of Parameters in the Configuration File.

a. In */etc/cinder*, use *vi* command to create a Huawei-customized driver configuration file names "*/etc/cinder/cinder_huawei_conf.xml*".

b. Configure parameters in the driver configuration file. You must change the bold parameters in this configuration files according to your production environment, such as RestURL, account, LUNType, StoragePool and Cinder node HostIP etc.

```
<?xml version='1.0' encoding='UTF-8'?>
<config>
        <Storage>
            <Product>V3</Product>
            <Protocol>FC</Protocol>
            <RestURL>https://100.115.10.224:8088/deviceManager/rest/</RestURL>
            <UserName>admin</UserName>
            <UserPassword>Admin@storage1</UserPassword>
        </Storage>
        <LUN>
                <LUNType>Thick</LUNType>
                <WriteType>1</WriteType>
                <LUNcopyWaitInterval>5</LUNcopyWaitInterval>
                <Timeout>432000</Timeout>
                <StoragePool>StoragePoolName</StoragePool>
        </LUN>
        <Host HostIP="192.168.8.6" HostType="Linux" />
</config>
```

Important notes:

- For more details about the parameters in the configuration file, please check below Parameters in the Configuration File. You can also find the details in section *4.2.7 Parameters in the Configuration File* at below link: https://github.com/huaweistorage/OpenStack_Driver/blob/master/ReleaseDoc/en/OpenStack%20Cinder%20Driver%20Configuration%20Guide.pdf

Table 1 Mandatory parameters

Parameter	Default Value	Description	Applicable To
Product	-	Type of a storage product. Such as **V3**.	All

Parameter	Default Value	Description	Applicable To
Protocol	-	Type of a connection protocol. Possible value is **FC**.	All
RestURL	-	Access address of the REST interface, **https://x.x.x.x:8088/devicemanager/rest/**. **x.x.x.x** is the management IP address.	V3
UserName	-	Storage administrator username.	All
UserPassword	-	Storage Password of an administrator.	All
StoragePool	-	Name of a storage pool to be used.	All
OSType	Linux	Operating system of the Nova compute node's host.	All
HostIP	-	IP address of the Nova compute node's host.	All

Note:
For the configuration of other products like OceanStor Dorado V3 series, TV2 storage and OceanStor 18000 high-end storage or other protocol like iSCSI, please also reference:
https://github.com/huaweistorage/OpenStack_Driver/blob/master/ReleaseDoc/en/OpenStack%20Cinder%20Driver%20Configuration%20Guide.pdf

Table 2 Optional parameters

Parameter	Default Value	Description	Applicable To
LUNType	Thick	Type of the LUNs to be created. The value can be **Thick** or **Thin**.	V3
WriteType	1	Cache write type. Possible values are: **1** (write back), **2** (write through), and **3** (mandatory write back).	All
LUNcopyWaitInterval	5	After LUN copy is enabled, the plug-in frequently queries the copy progress. You can set a value to specify the query interval.	All
Timeout	432000	Timeout interval for waiting LUN copy of a storage device to complete. The unit is second.	V3
OSType	Linux	Operating system of the Nova compute node's host.	All
HostIP	-	IP address of the Nova compute node's host.	All

Notes:
You need to get Huawei's license first before creating thin lun.

c. Change *cinder_huawei_conf.xml* file owner to cinder user and cinder user group to allow cinder service to full access permission.

root@node-6:~# chown cinder:cinder /etc/cinder/cinder_huawei_conf.xml

d. Configure the cinder.conf file

Backup the*/etc/cinder/cinder.conf* file before any modify. In the *[default]* block of */etc/cinder/cinder.conf*, add the following contents. *enabled_backends* indicates the default backend will be used when creating volumes. *volume_driver* indicates the loaded driver file, and *cinder_huawei_conf_file* indicates the specified Huawei-customized configuration file.

enabled_backends=OceanStor_V3

[OceanStor_V3]
volume_driver = cinder.volume.drivers.huawei.huawei_driver.HuaweiV3FCDriver
cinder_huawei_conf_file = /etc/cinder/cinder_huawei_conf.xml
volume_backend_name = OceanStor_V3

e. Run the service cinder-volume restart command to restart the Block Storage service.

root@node-6:~# service cinder-volume restart

f. Once the *cinder.conf* and Huawei customized driver configuration files have been created, cinder commands need to be issued in order to create and associate OpenStack volume types with the declared volume_backend_names:

root@node-6:~# cinder type-create OceanStor_V3
root@node-6:~# cinder type-key OceanStor_V3 set volume_backend_name=OceanStor_V3
root@node-6:~# cinder extra-specs-list

```
[root@node-7 ~(keystone_admin)]# cinder  extra-specs-list
+--------------------------------------+-------------+---------------------------------------+
|                  ID                  |     Name    |               extra_specs             |
+--------------------------------------+-------------+---------------------------------------+
| 9a9d1f0a-2451-4774-a635-745f208e2980 | OceanStor_V3 | {u'volume_backend_name': u'OceanStor_V3'} |
+--------------------------------------+-------------+---------------------------------------+
```

g. Check cinder *OceanStor_V3* volume type on OpenStack Dashboard:

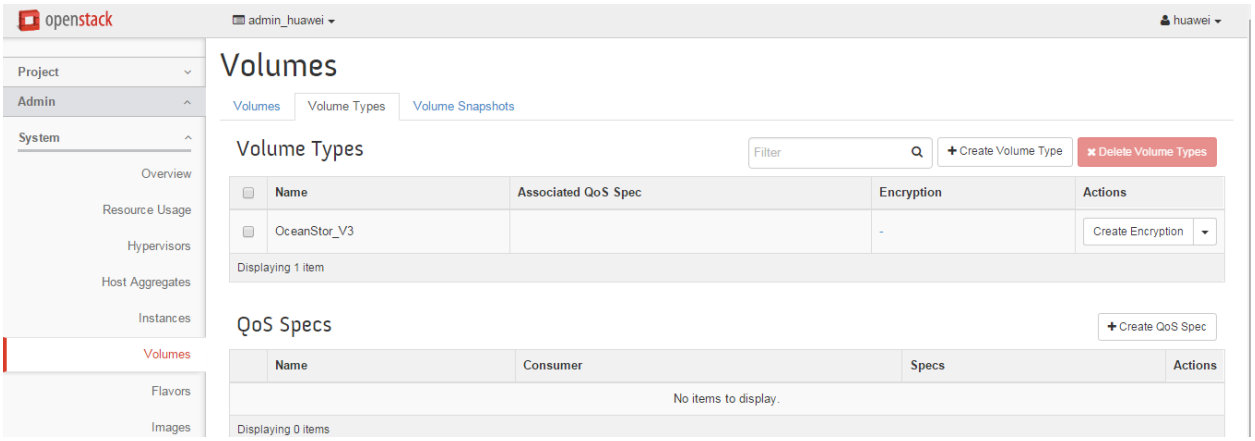

6. Testing

6.1 Test tools

Equipment	Purpose
Fuel health check scripts	Functional/acceptance testing
other	Test Huawei Volume Driver Function

6.2 Test cases

- Fuel Health Check test cases, see Fuel Health Check test result 6.3.1
- Other manually test cases:

Function	Description	Remark
Create Volume	Create OceanStor_V3 volume type empty volume on OpenStack Dashboard or CLI	pass
Show Volume	Query OceanStor_V3 volume type volume on OpenStack Dashboard or CLI	pass
Delete Volume	Delete OceanStor_V3 volume type volume on OpenStack Dashboard or CLI	pass
Create Snapshot	Create OceanStor_V3 volume type volume snapshot on OpenStack Dashboard or CLI	pass
Show Snapshot	Query OceanStor_V3 volume type volume snapshot on OpenStack Dashboard or CLI	pass
Delete Snapshot	Delete OceanStor_V3 volume type volume snapshot on OpenStack Dashboard or CLI	pass
Create Volume From Image	Create OceanStor_V3 volume type volume from image on OpenStack Dashboard or CLI	pass
Create Volume From Volume	Create OceanStor_V3 volume type volume from volume on OpenStack Dashboard or CLI	pass

Create Volume From Snapshot	Create OceanStor_V3 volume type volume from snapshot on OpenStack Dashboard or CLI	pass
Launch Instance From Volume	Launch Instance From Volume on OpenStack Dashboard or CLI	pass
Attach Volume to Instance	Attach Volume to Instance on OpenStack Dashboard or CLI	pass
Detach Volume to Instance	Detach Volume to Instance on OpenStack Dashboard or CLI	pass

6.2.1 Target Use case(s)

1) Large-scale cloud computing data centers

Huawei Volume Driver enables Huawei Storages integrate with MOS, providing converged functionality that maximizes performance and capacity utilization, Huawei storage products meet a wide range of needs: entry-level, mid-range, and high-end enterprise applications; Big Data; vertical industry applications; and cloud storage applications.

2) Critical enterprise IT infrastructure

MOS integrate with Huawei Storages, providing the best data services for enterprise, mission-critical businesses.

6.2.2 Deployment modes and configuration options

OS	Mode	HV	Network	Storage
			Neutron with VLAN	Cinder driver
Ubuntu	HA	KVM	x	x

6.3 Test results (if FUEL HealthCheck is used)

6.3.1 Fuel Health Check Result

The Fuel Health Check result is tested with Huawei Volume Driver and Huawei OceanStor V3 storage.

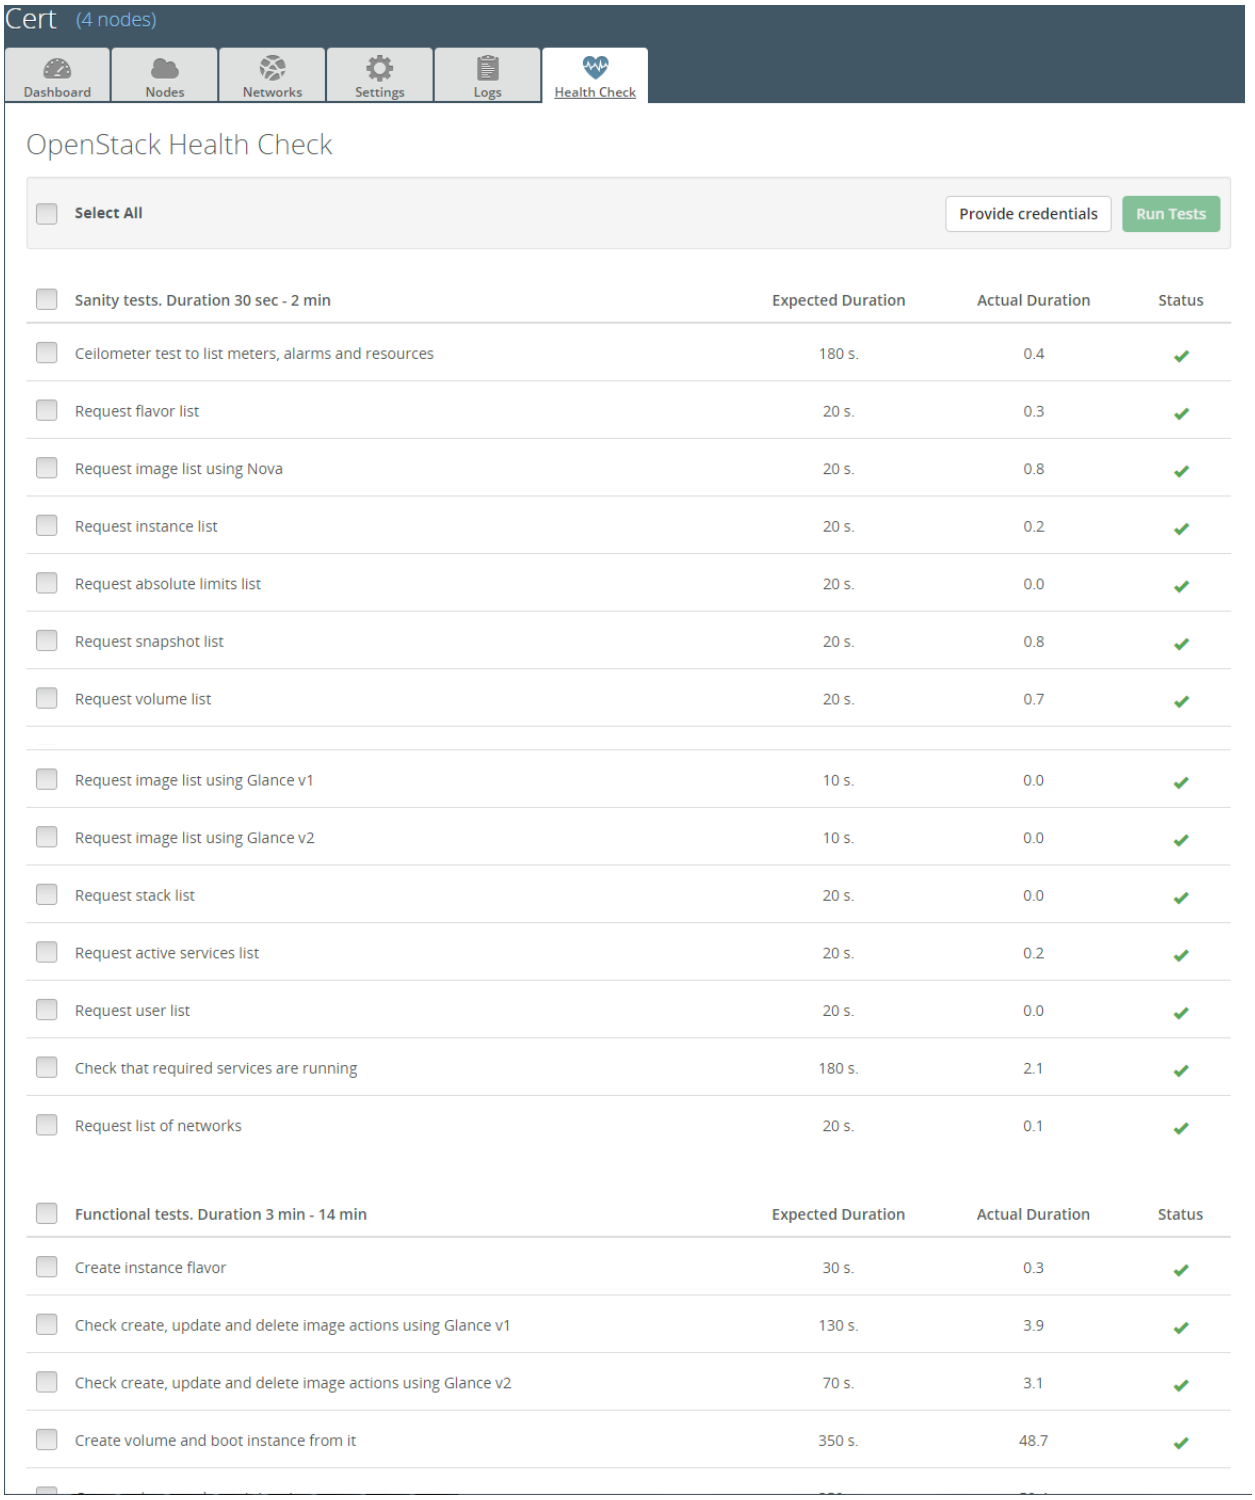
Cert (4 nodes)
Dashboard
Nodes
Networks
Settings
Logs
Health Check
OpenStack Health Check
Select All
Provide credentials
Run Tests
Sanity tests. Duration 30 sec - 2 min
Expected Duration
Actual Duration
Status
Ceilometer test to list meters, alarms and resources
180 s.
0.4
Request flavor list
20 s.
0.3
Request image list using Nova
20 s.
0.8
Request instance list
20 s.
0.2
Request absolute limits list
20 s.
0.0
Request snapshot list
20 s.
0.8
Request volume list
20 s.
0.7
Request image list using Glance v1
10 s.
0.0
Request image list using Glance v2
10 s.
0.0
Request stack list
20 s.
0.0
Request active services list
20 s.
0.2
Request user list
20 s.
0.0
Check that required services are running
180 s.
2.1
Request list of networks
20 s.
0.1
Functional tests. Duration 3 min - 14 min
Expected Duration
Actual Duration
Status
Create instance flavor
30 s.
0.3
Check create, update and delete image actions using Glance v1
130 s.
3.9
Check create, update and delete image actions using Glance v2
70 s.
3.1
Create volume and boot instance from it
350 s.
48.7

Create volume and attach it to instance	350 s.	59.4	✔
Check network connectivity from instance via floating IP	300 s.	96.6	✔
Create keypair	25 s.	0.5	✔
Create security group	25 s.	0.5	✔
Check network parameters	50 s.	0.1	✔
Launch instance	200 s.	23.0	✔
Launch instance with file injection	200 s.	26.1	✔
Launch instance, create snapshot, launch instance from snapshot	300 s.	59.7	✔
Create user and authenticate with it to Horizon	80 s.	0.3	✔

HA tests. Duration 30 sec - 8 min	**Expected Duration**	**Actual Duration**	**Status**
Check data replication over mysql	10 s.	30.8	✔
Check if amount of tables in databases is the same on each node	10 s.	3.8	✔
Check galera environment state	10 s.	1.6	✔
Check pacemaker status	10 s.	1.7	✔
RabbitMQ availability	100 s.	13.5	✔
RabbitMQ replication	100 s.	25.3	✔

Platform services functional tests. Duration 3 min - 60 min	**Expected Duration**	**Actual Duration**	**Status**
Ceilometer test to check alarm state and get Nova metrics	60 s.	80.5	✔
Ceilometer test to check notifications from Glance	5 s.	3.0	✔
Ceilometer test to check notifications from Keystone	5 s.	2.6	✔
Ceilometer test to check notifications from Neutron	40 s.	3.2	✔
Ceilometer test to check notifications from Cinder	10 s.	3.4	✔
Ceilometer test to create, check and list samples	5 s.	0.5	✔
Ceilometer test to create, update, check and delete alarm	120 s.	120.5	✔
Typical stack actions: create, delete, show details, etc.	560 s.	27.6	✔

Check pacemaker status	10 s.	1.7	✔
RabbitMQ availability	100 s.	13.5	✔
RabbitMQ replication	100 s.	25.3	✔

Platform services functional tests. Duration 3 min - 60 min	Expected Duration	Actual Duration	Status
Ceilometer test to check alarm state and get Nova metrics	60 s.	80.5	✔
Ceilometer test to check notifications from Glance	5 s.	3.0	✔
Ceilometer test to check notifications from Keystone	5 s.	2.6	✔
Ceilometer test to check notifications from Neutron	40 s.	3.2	✔
Ceilometer test to check notifications from Cinder	10 s.	3.4	✔
Ceilometer test to create, check and list samples	5 s.	0.5	✔
Ceilometer test to create, update, check and delete alarm	120 s.	120.5	✔
Typical stack actions: create, delete, show details, etc.	560 s.	27.6	✔
Advanced stack actions: suspend, resume and check	660 s.	66.7	✔
Check stack autoscaling	2200 s.	242.1	✔
Check stack rollback	310 s.	105.2	✔
Update stack actions: inplace, replace and update whole template	950 s.	90.9	✔

Cloud validation tests. Duration 30 sec - 2 min	Expected Duration	Actual Duration	Status
Check disk space outage on controller and compute nodes	20 s.	1.7	✔
Check log rotation configuration on all nodes	20 s.	1.8	✔

Configuration tests. Duration 30 sec - 2 min	Expected Duration	Actual Duration	Status
Check usage of default credentials on master node	20 s.	38.4	✔
Check if default credentials for OpenStack cluster have changed	20 s.	0.0	✔
Check usage of default credentials for keystone on master node	20 s.	0.2	✔

6.3.2 Other manually testing result screenshots.

We want to check that all the volume operations for Huawei OceanStor_V3 volume type are working normally. The volumes used in following test cases are all created in OceanStor_V3 volume type.

a. Created cirros image list on Openstack Dashboard. We use the default cirros image which created by installing OpenStack cloud environment for creating volume from image test case.

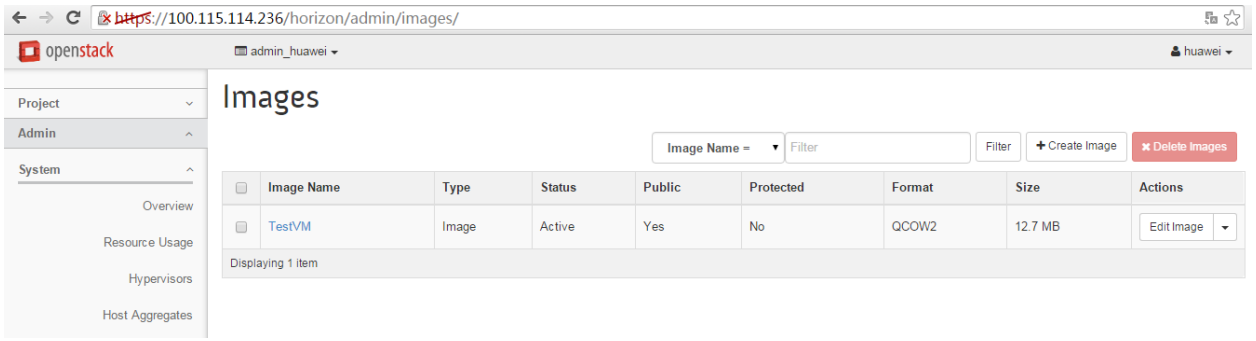

b. Volume operations: Created OceanStor_V3 type volume list on Openstack Dashboard.

- Create empty volume: create a empty volume from OceanStor_V3 volume type.
- Create volume from image: create a volume from cirros image.
- Create volume from snapshot: take a snapshot for a specific volume and create a volume from the snapshot.
- Create volume from volume: create a volume from an OceanStor_V3 volume.
- Delete: delete an OceanStor_V3 type volume.
- Modify: modify the size of a OceanStor_V3 type volume.

The volume list created in OceanStor_V3 volume type.

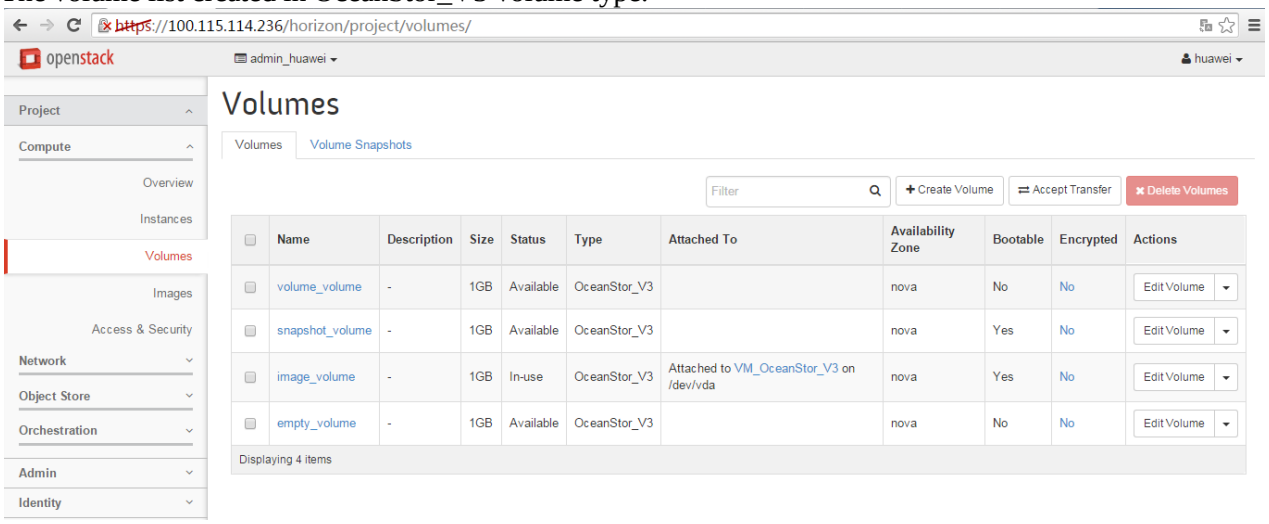

The volume snapshots list

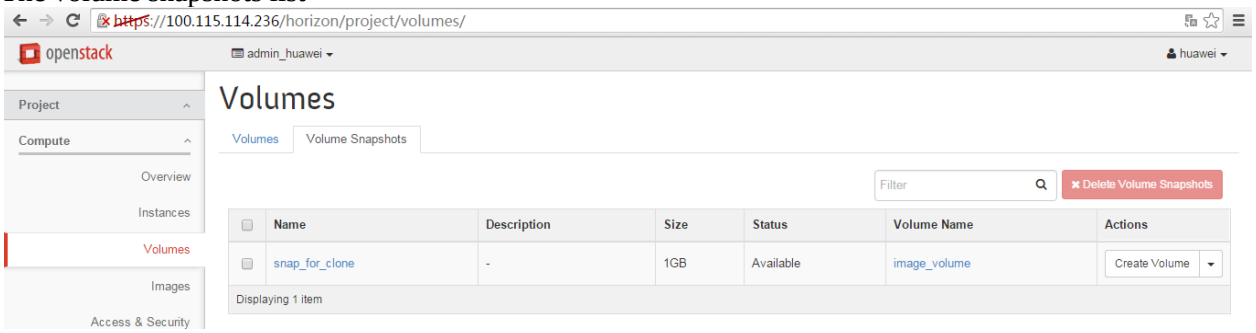

c. Created Volume Snapshots list on Openstack Dashboard.

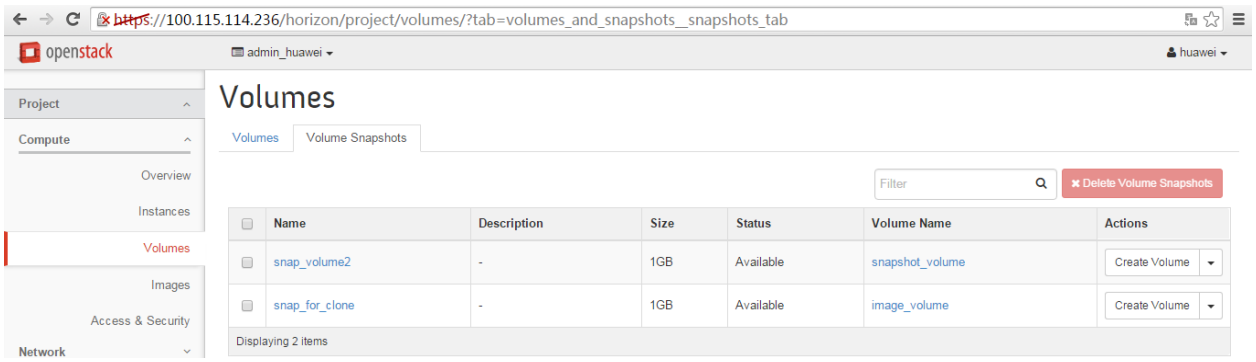

d. Created Instances with OceanStor_V3 volumes on Openstack dashboard. (Also checked delete, modify, etc)

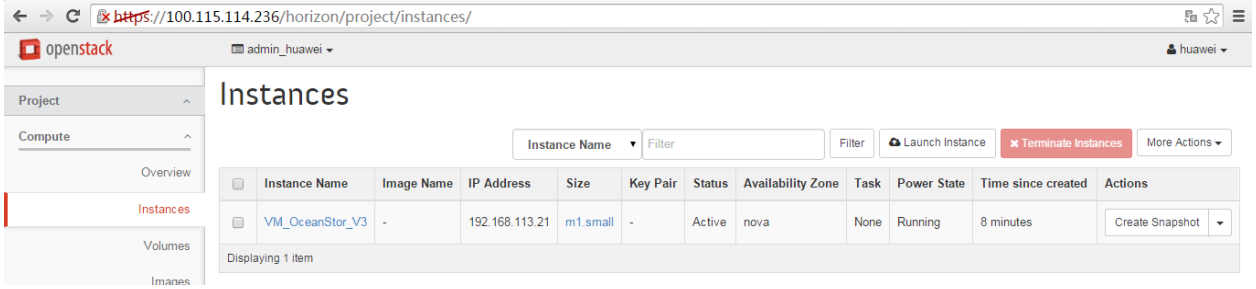

VM instance console screenshot on Openstack dashboard.

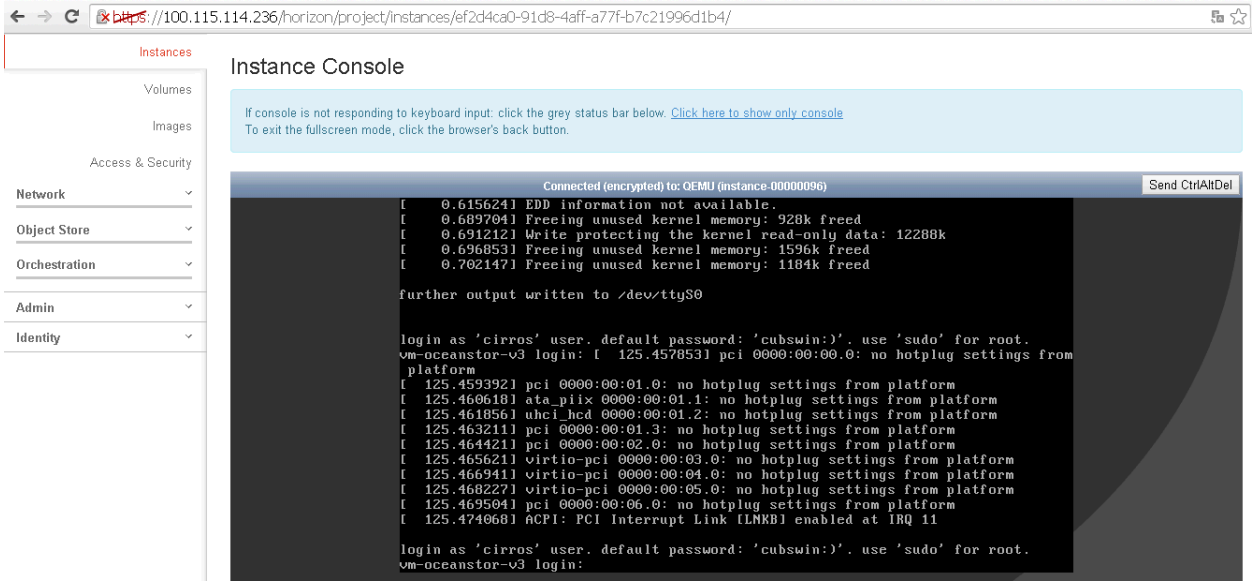

e. Attach/detach an OceanStor_V3 volume to instance on Openstack Dashboard

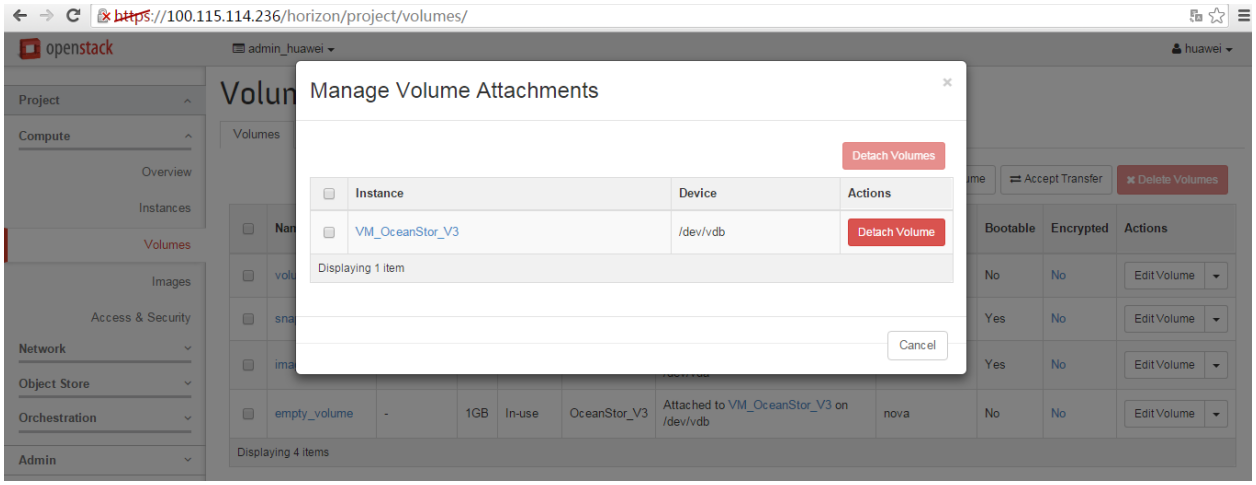

7. Appendix.

7.1 How to download Huawei Volume Driver software.

Visiting https://github.com/huaweistorage/OpenStack_Driver/tree/master

7.2 How to get the Huawei Volume Driver configuration guide

https://github.com/huaweistorage/OpenStack_Driver/blob/master/ReleaseDoc/en/OpenStack%20Cinder%20Driver%20Configuration%20Guide.pdf

7.3 How to get the Huawei Storage document

Visiting http://e.huawei.com/en/products/cloud-computing-dc/storage