Installing the Cisco Nexus 1000V

This chapter contains the following sections:

- Prerequisites, page 1
- Guidelines and Limitations, page 1
- Planning for Cisco Nexus 1000V Switch Installation , page 2
- Installing the Undercloud, page 2
- Installing the Overcloud, page 3
- Managing Graceful Controller Restart, page 9
- Basic Troubleshooting , page 11

Prerequisites

Before you begin the Undercloud and Overcloud installation procedure, ensure that you have:

- Downloaded and installed RHEL 7.2 server image to the server where you intend to deploy the Undercloud. For more information about how to download and install RHEL 7.2 server image, see Red Hat Enterprise Linux Product Download.
- Downloaded deployment Ramdisk for RHEL-OSP Director 7.2 from RedHat repository.
- Downloaded discovery Ramdisk for RHEL-OSP Director 7.2 from RedHat repository.
- Downloaded Overcloud image for RHEL-OSP Director 7.2 from RedHat repository.
- Installed the libguestfs-tools package from RedHat repository. This package is required for customization commands, such as **virt-customize**.

Guidelines and Limitations

- In VXLAN mulicast mode, the VTEP ports on VEM need to respond to incoming IGMP query traffic for the multicast group to which they belong. However, the default firewall rules (iptables) drop the

incoming IGMP query traffic from reaching the VTEP interfaces. In order to allow this traffic, a firewall rule needs to be configured on the respective compute and network hosts, as follows:

```
#iptables -I INPUT 1 -p igmp -j ACCEPT
```

- If no other interfaces other than the management interface come up after you reboot, you can either manually bring up the interfaces by entering **ifconfig** *interface_name* or change the ONBOOT parameter to **yes** in the /etc/sysconfig/network and the /etc/sysconfig/network-scripts/ifcfg-"interface name" files.

Planning for Cisco Nexus 1000V Switch Installation

The Cisco Nexus 1000V Switch installation using Red Hat Enterprise Limited OpenStack Platform involves setting up two OpenStack environments, Undercloud and Overcloud. The single-system OpenStack environment that is used to configure and provision the target OpenStack environment is known as the Undercloud. The resulting OpenStack environment on which you intend to deploy the applications is known as Overcloud. The installation procedure involves planning the Undercloud and Overcloud environment.

- Planning for Undercloud involves creating a non-root user and creating directories for system images and Heat templates. For detailed information about Undercloud installation and planning, see Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage.
- Planning for Overcloud involves planning for different types of nodes required, network topology, and storage options in the cloud environment. For detailed information about Overcloud installation and planning, see Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage.

Installing the Undercloud

The Undercloud is the main director node that provides components to configure and manage the OpenStack nodes in the Overcloud. Installing the Undercloud involves:

- Creating a Non-Root User for Director Installation
- Creating Directories for System Images and Heat Templates
- Configuring IP Forwarding support on the Director Host
- Setting the Host name for the Director Host
- Registering the Director Host
- Installing the Director
- Configuring the Director
- Obtaining Images for the Overcloud Nodes
- Configuring DNS for the Overcloud Nodes

Before you proceed to Overcloud installation, ensure that the Undercloud is successfully installed and available as described in the RHEL-OSP documentation available at: Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage.

Installing the Overcloud

Complete these steps to deploy an Overcloud environment. For detailed Overcloud installation information, see the RHEL-OSP documentation available at: Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage.

Before You Begin

Ensure that the Undercloud is installed and running.

Procedure

Step 1 Log into the Undercloud OpenStack platform as a non-root user. For example:

```
#ssh root@<undercloud-machine-ip>

[root@undercloud ~]#su - stack
```

Step 2 Source the environment variables using the **source** command. For example:

```
[stack@undercloud ~]$source stackrc
```

Step 3 Add information about Overcloud nodes in the configuration file, instackenv.json. If the configuration file does not exist, you can create a new configuration file. For more information about the configuration file, see the RHEL-OSP documentation available at: Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage.

Step 4 Copy the setup image files from the official Red Hat repository to the Undercloud platform. See the Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage for exact location of the image files.

Step 5 Extract the setup image files using the **tar -xvf** command. For example:

```
[stack@undercloud ~]$tar -xvf SetupImage.tar
```

Step 6 Customize the extracted setup files with VSM and VEM information. You can also download the sample script file, available at https://cnsg-yum-server.cisco.com/yumrepo/osp7/n1kv-injection.sh. The sample script or the set of commands below will need to have the Red Hat subscription information populated. For example:

```
# Enable subscription manager
virt-customize -a overcloud-full.qcow2 --run-command 'subscription-manager register
--username=YourRedHatSubscriptionName --password=YourRedHatSubscriptionPassword'
virt-customize -a overcloud-full.qcow2 --run-command 'subscription-manager attach --pool
YourPoolId'

# Create repo file for hosted repo
echo "[n1kv]
name=n1kv
baseurl=https://cnsg-yum-server.cisco.com/yumrepo/
enabled=1
gpgcheck=0" > n1kv.repo

# Add hosted repo to image
virt-customize -a overcloud-full.qcow2 --upload n1kv.repo:/etc/yum.repos.d/

# Cleanup repo file
rm n1kv.repo
```

```
# Install VEM
virt-customize -a overcloud-full.qcow2 --install nexus1000v

# Install VSM
virt-customize -a overcloud-full.qcow2 --install nexus-1000v-iso

# Unregister the node
virt-customize -a overcloud-full.qcow2 --run-command 'subscription-manager remove --all'
virt-customize -a overcloud-full.qcow2 --run-command 'subscription-manager unregister'
```

Step 7 Upload the image files to the Undercloud using the **openstack overcloud image upload** command. For example:

```
[stack@undercloud ~]$openstack overcloud image upload
```

Step 8 Verify whether the image files are successfully uploaded to the Undercloud using the **glance** command. For example:

```
stack@undercloud ~]$ glance image-list
+--------------------------------------+------------------------+------------+---------------+----------+-------+
| ID                                   | Name                   | Disk Format | Container
Format | Size       | Status |
+--------------------------------------+------------------------+------------+---------------+----------+-------+
| 735c6856-ba6b-4d33-962f-db168cd3078d | bm-deploy-kernel       | aki         | aki
       | 5027584    | active |
| 67c16270-cfe8-4979-89bf-a840233cce95 | bm-deploy-ramdisk      | ari         | ari
       | 56302611   | active |
| 3a21da43-f4fa-4629-84b9-df8d91543325 | overcloud-full         | qcow2       | bare
       | 1342570496 | active |
| 81e5f8d4-20e3-4192-b65e-396faeb64b9e | overcloud-full-initrd  | ari         | ari
       | 36757801   | active |
| 02cd5efc-2a75-4cb9-95ea-e7984aeff359 | overcloud-full-vmlinuz | aki         | aki
       | 5027584    | active |
+--------------------------------------+------------------------+------------+---------------+----------+-------+
```

Step 9 Edit the YAML file (cisco-n1kv-config.yaml) to configure VSM and VEM. For example:

```
[root@undercloud ~]#vi
/usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml
# A Heat environment file which can be used to enable a
# a Cisco N1KV backend, configured via puppet
resource_registry:
  OS::TripleO::ControllerExtraConfigPre:
../puppet/extraconfig/pre_deploy/controller/cisco-n1kv.yaml
  OS::TripleO::ComputeExtraConfigPre:
../puppet/extraconfig/pre_deploy/controller/cisco-n1kv.yaml
  OS::TripleO::Controller: ../puppet/controller-puppet.yaml

parameter_defaults:
  N1000vVSMIP: '16.0.0.12'
  N1000vMgmtGatewayIP: '16.0.0.1'
  N1000vVSMDomainID: '235'
  N1000vVSMHostMgmtIntf: br-ex
  N1000vPacemakerControl: true
  N1000vExistingBridge: true
  N1000vVSMVersion: '5.2.1.SK3.2.2b-1'
  N1000vVEMHostMgmtIntf: 'br-ex'
  N1000vUplinkProfile: '{eth2: sys-uplink,}
  NeutronServicePlugins:
```

```
"router,networking_cisco.plugins.ml2.drivers.cisco.n1kv.policy_profile_service.PolicyProfilePlugin"

  NeutronTypeDrivers: "vlan,vxlan"

  NeutronCorePlugin: "neutron.plugins.ml2.plugin.Ml2Plugin"
NodeDataLookup: |
{
"3a813b2b-c591-44c0-bc48-7b6ef27429e0":    # This is the System UUID.
{
"neutron::agents::n1kv_vem::uplink_profile":{"eth1": "system-uplink-macpin",
"eth2":"system-uplink-macpin"},
# Specify the Uplink port-profiles to be used for respective interfaces
"neutron::agents::n1kv_vem::vtep_config":{"vtep1":{"profile": "vtep1-pp", "ipmode":"static",
 "ipaddress": "15.51.0.11", "netmask": "255.255.255.0"}},
# Specify the vtep name (in case you need to create a vtep) and associated port-profile to
 be used. Provide the mode as static or dhcp as per your network deployment. If static is
chosen, then provide the IP address and Netmask.
"neutron::agents::n1kv_vem::host_mgmt_intf": "enp8s0"
# Specify the host's management interface using which VEM will communicate with VSM over
Layer 3.
}
}
```

Note After you deploy the VSM, ensure that you define the uplink port profile and VTEP port profile in the VSM using the same names as specified in the configuration file.

Note For more information about Heat parameters, see Heat Templates.

Note An Overcloud deployed using this image will not switch to Openvswitch for networking.

Step 10 Register the Overcloud nodes with Openstack Ironic service on Undercloud using the **openstack baremetal import** command. For example:

```
[stack@undercloud ~]$openstack baremetal import --json instackenv.json
```

Step 11 Assign kernel and Ramdisk to nodes using the **openstack baremetal configure boot** command. For example:

```
[stack@undercloud ~]$openstack baremetal configure boot
```

Step 12 Verify the registered nodes using the **ironic node-list** command. For example:

```
[stack@undercloud ~]$ ironic node-list
+--------------------------------------+------+--------------+-------------+-----------------+-------------+
| UUID                                 | Name | Instance UUID | Power State | Provision
State | Maintenance |
+--------------------------------------+------+--------------+-------------+-----------------+-------------+
| e2122ad2-4953-4528-ac9f-ef2355e5adee | None | None          | power on    | manageable
    | False       |
| 8ace9caf-bb69-44e8-a202-ed17bec074b3 | None | None          | power on    | manageable
    | False       |
+--------------------------------------+------+--------------+-------------+-----------------+-------------+
```

Note Provision state can be either available or manageable. You can use the **openstack baremetal introspection bulk start** command, which changes the Provision State to manageable if it is in available state before the inspection starts.

Step 13 Introspect the hardware attributes of the Overcloud nodes using the **openstack baremetal introspection bulk start** command. For example:

```
[stack@undercloud ~]$ openstack baremetal introspection bulk start
Setting available nodes to manageable...
Starting introspection of node: e2122ad2-4953-4528-ac9f-ef2355e5adee
Starting introspection of node: 8ace9caf-bb69-44e8-a202-ed17bec074b3
```

```
Waiting for discovery to finish...
Discovery for UUID 8ace9caf-bb69-44e8-a202-ed17bec074b3 finished successfully.
Discovery for UUID e2122ad2-4953-4528-ac9f-ef2355e5adee finished successfully.
Setting manageable nodes to available...
Node e2122ad2-4953-4528-ac9f-ef2355e5adee has been set to available.
Node 8ace9caf-bb69-44e8-a202-ed17bec074b3 has been set to available.
Discovery completed.

[stack@undercloud ~]$ ironic node-list
+--------------------------------------+------+--------------+-------------+-----------------+-------------+
| UUID                                 | Name | Instance UUID | Power State | Provision
State | Maintenance |
+--------------------------------------+------+--------------+-------------+-----------------+-------------+
| e2122ad2-4953-4528-ac9f-ef2355e5adee | None | None          | power off   | available
    | False       |
| 8ace9caf-bb69-44e8-a202-ed17bec074b3 | None | None          | power off   | available
    | False       |
+--------------------------------------+------+--------------+-------------+-----------------+-------------+
```

Step 14 Create the flavor for compute node, controller node, and baremetal kernel image using the **openstack flavor create** command. For example:

```
$ openstack flavor create --id auto --ram 4096 --disk 40 --vcpus 1 baremetal
$ openstack flavor set --property "cpu_arch"="x86_64" –property
"capabilities:boot_option"="local" --property "capabilities:profile"="baremetal" baremetal

$ openstack flavor create --id auto --ram 8192 --disk 80 --vcpus 4 control
$ openstack flavor set --property "cpu_arch"="x86_64" –property
"capabilities:boot_option"="local" --property "capabilities:profile"="control" control

$ openstack flavor create --id auto --ram 8192 --disk 80 --vcpus 4 compute
$ openstack flavor set --property "cpu_arch"="x86_64" –property
"capabilities:boot_option"="local" --property "capabilities:profile"="compute" compute
```

Note You need to tag your nodes with the hardware profiles manually or using Automated Health Check (AHC) functionality. For detailed information about manual tagging of nodes, see Manually Tagging the Nodes. For detailed information about AHC, see Automatically Tagging Nodes with Automated Health Check (AHC) Tools.

Step 15 List the nodes using **nova flavor-list** command. For example:

```
[stack@undercloud ~]$ nova flavor-list
+--------------------------------------+---------+---------+-----+---------+-----+-------+-----------+---------+
| ID                                   | Name      | Memory_MB | Disk | Ephemeral | Swap |
 VCPUs | RXTX_Factor | Is_Public |
+--------------------------------------+---------+---------+-----+---------+-----+-------+-----------+---------+
| 22446161-c471-42fb-9b48-b8944e961393 | control   | 8192      | 80   | 0         |      |
 4     | 1.0         | True      |
| 4460dd96-63be-4c40-9082-7c1093400d8c | compute   | 8192      | 80   | 0         |      |
 4     | 1.0         | True      |
| f5717ee5-87cc-4d18-8627-4c3a6d57754f | baremetal | 4096      | 40   | 0         |      |
 1     | 1.0         | True      |
+--------------------------------------+---------+---------+-----+---------+-----+-------+-----------+---------+
```

Step 16 Deploy the Overcloud using the **openstack overcloud deploy** command. Include the Cisco Nexus 1000V switch configuration file (the file, cisco-n1kv-config.yaml, that you configured in the previous step) in the

command to ensure that the VEM and VSM are deployed with the configuration specified in the configuration file. For example:

```
[stack@undercloud ~]$ openstack overcloud deploy --templates
--ceph-storage-scale 0
--control-scale 3
--control-flavor control
--compute-scale 3
--compute-flavor compute
-e /usr/share/openstack-tripleo-heat-templates/overcloud-resource-registry-puppet.yaml
-e /usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml
--neutron-network-type vlan
--neutron-network-vlan-ranges datacentre:1551:1600
--neutron-tunnel-types vlan
--swift-storage-scale 0
--swift-storage-flavor compute
--ceph-storage-flavor compute
--block-storage-flavor compute
--ntp-server ntp.esl.cisco.com

[root@undercloud ~]#
```

Note The configuration file also leverages two parameters, OS::TripleO::ControllerExtraConfigPre and OS::TripleO::ComputeExtraConfigPre. You can use these two parameters to integrate third-party components and provide additional Heat configuration. Each parameter allows you to configure multiple components.

Step 17 Verify the deployment using the **nova list** command. For example:

```
[stack@undercloud ~]$ nova list
+--------------------------------------+-------------------------+--------+------------+-------------+------------------+
| ID                                   | Name                    | Status | Task State |
Power State | Networks            |
+--------------------------------------+-------------------------+--------+------------+-------------+------------------+
| c6babe19-6424-421a-b081-fb8026885dbc | overcloud-compute-0     | ACTIVE | -          |
Running     | ctlplane=16.0.0.109 |
| 24336715-4a78-4c74-aede-bc70ce098c39 | overcloud-controller-0 | ACTIVE | -          |
Running     | ctlplane=16.0.0.110 |
+--------------------------------------+-------------------------+--------+------------+-------------+------------------+
```

Step 18 Log into the Overcloud Controller node using **SSH**. For example:

```
[stack@undercloud ~]$ssh heat-admin@16.0.0.110
The authenticity of host '16.0.0.110 (16.0.0.110)' can't be established.
ECDSA key fingerprint is e8:04:24:dc:2d:09:f2:c5:be:da:21:8d:72:41:47:70.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '16.0.0.110' (ECDSA) to the list of known hosts.

[heat-admin@overcloud-controller-0 ~]$
```

Step 19 Verify the value the following parameters in the neutron.conf file:

```
[stack@undercloud ~]$ ssh admin@16.0.0.12
Nexus 1000v Switch
Password:

Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2015, Cisco Systems, Inc. All rights reserved.
```

```
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
```

Note If the parameters values do not match to the values stated above, set the parameters to specified values and run the following commands in a sequence:

1. **service neutron-server restart**
2. **service neutron-l3-agent restart**
3. **service neutron-dhcp-agent restart**

Step 20 SSH into VSM using the VSM IP address configured in the YAML Heat template. For example:

Step 21 Verify the configured nodes using the **show module** command. For example:

```
vsm-p# show module
Mod  Ports  Module-Type                      Model               Status
---  -----  --------------------------------  ------------------  ------------
1    0      Virtual Supervisor Module         Nexus1000V          active *
2    0      Virtual Supervisor Module         Nexus1000V          ha-standby
3    1022   Virtual Ethernet Module           NA                  ok
4    1022   Virtual Ethernet Module           NA                  ok
5    1022   Virtual Ethernet Module           NA                  ok
6    1022   Virtual Ethernet Module           NA                  ok
7    1022   Virtual Ethernet Module           NA                  ok

Mod  Sw                 Hw
---  -----------------  ------------------------------------------------
1    5.2(1)SK3(2.2b)    0.0
2    5.2(1)SK3(2.2b)y    0.0
3    5.2(1)SK3(2.2b)    Linux 3.10.0-229.7.2.el7.x86_64
4    5.2(1)SK3(2.2b)    Linux 3.10.0-229.14.1.el7.x86_64
5    5.2(1)SK3(2.2b)    Linux 3.10.0-229.14.1.el7.x86_64
6    5.2(1)SK3(2.2b)    Linux 3.10.0-229.14.1.el7.x86_64
7    5.2(1)SK3(2.2b)    Linux 3.10.0-229.7.2.el7.x86_64

Mod  Server-IP        Server-UUID                          Server-Name
---  ---------------  -----------------------------------  --------------------
1    16.0.0.12        NA                                   NA
2    16.0.0.12        NA                                   NA
3    16.0.0.189       3858C23E-1681-E411-0000-00000000000D  overcloud-compute-1.localdomain
4    16.0.0.188       3858C23E-1681-E411-0000-00000000000C  overcloud-controller-1.localdomain
5    16.0.0.191       3858C23E-1681-E411-0000-000000000009  overcloud-controller-2.localdomain
6    16.0.0.192       3858C23E-1681-E411-0000-00000000000E  overcloud-controller-0.localdomain
7    16.0.0.190       3858C23E-1681-E411-0000-00000000000B  overcloud-compute-0.localdomain

* this terminal session
vsm-p#
```

Step 22 Configure the default port profile and uplink profile for different types of ports such as uplink ports. The default profile should have basic **no shutdown** configuration because it is used for router and DHCP port.

The uplink port should be configured as a trunk and optionally can be configured with a range of allowed VLANs (if the VLAN range is not specified, all the VLANs are allowed). Save this configuration to startup using the **copy** command. For example:

```
[stack@undercloud ~]$ ssh admin@16.0.0.12
conf t
port-profile default-pp
no shut
state enabled
publish port-profile
port-profile type ethernet system-uplink
switchport mode trunk
no shut
state enabled
publish port-profile
end
copy r s
```

Step 23 After you complete the Overcloud configuration, backup the static configuration from VSM and copy it to a remote location. For example:

```
vsm-p# show running-config static > sftp://stack@192.0.2.1/home/stack/backup.txt
```

Note Backup copy is helpful in recovering VSM state in the event of multiple controller failure.

Managing Graceful Controller Restart

Complete these steps to ensure graceful controller restart.

Procedure

Step 1 Identify the primary and secondary VSM controller hosts using the **pcs status** command. For example:

```
[root@overcloud-controller-2 heat-admin]# pcs status | grep vsm
vsm-p (ocf::heartbeat:VirtualDomain): Started overcloud-controller-1
vsm-s (ocf::heartbeat:VirtualDomain): Started overcloud-controller-2
```

Step 2 Copy the primary VSM disk to other controller nodes.

a) Copy the primary VSM disk to the Undercloud. For example:

```
[stack@n1kv-ospd-test ~]$ nova list
+--------------------------------------+------------------------+--------+------------+-------------+-------------------+
| ID | Name | Status | Task State | Power State | Networks |
+--------------------------------------+------------------------+--------+------------+-------------+-------------------+
| 0206d790-6b56-4b3e-a589-51e9150c39b1 | overcloud-compute-0 | ACTIVE | - | Running |
ctlplane=192.0.2.9 |
| c0140cc1-9f5f-4f34-a124-8a78f31586db | overcloud-controller-0 | ACTIVE | - | Running
| ctlplane=192.0.2.10 |
| ff0e325e-bfa7-4af2-836b-3acd12f72670 | overcloud-controller-1 | ACTIVE | - | Running
| ctlplane=192.0.2.11 |
| 364c507b-3c79-407e-86bd-f3d1eea5aa1f | overcloud-controller-2 | ACTIVE | - | Running
| ctlplane=192.0.2.12 |
+--------------------------------------+------------------------+--------+------------+-------------+-------------------+
```

```
[stack@n1kv-ospd-test ~]$ ssh heat-admin@192.0.2.11
[heat-admin@overcloud-controller-1 ~]$ sudo su
[root@overcloud-controller-1 heat-admin]# scp /var/spool/cisco/vsm/primary_disk
stack@192.0.2.1:/home/stack/
```

b) Copy the primary VSM disk from Undercloud to a temporary location on the remaining nodes. For example:

```
[stack@n1kv-ospd-test ~]$ scp primary_disk heat-admin@192.0.2.10:/tmp
[stack@n1kv-ospd-test ~]$ scp primary_disk heat-admin@192.0.2.12:/tmp
```

c) Log in to remaining nodes with root privileges and copy disk file from temporary location to actual disk location. For example:

```
[stack@n1kv-ospd-test ~]$ ssh heat-admin@192.0.2.10
[heat-admin@overcloud-controller-0 ~]$ sudo su
[root@overcloud-controller-0 heat-admin]# cp /tmp/primary_disk
/var/spool/cisco/vsm/primary_disk
cp: overwrite /var/spool/cisco/vsm/primary_disk?
[root@overcloud-controller-0 heat-admin]# exit
[heat-admin@overcloud-controller-0 ~]$ exit
[stack@n1kv-ospd-test ~]$ ssh heat-admin@192.0.2.12
[heat-admin@overcloud-controller-0 ~]$ sudo su
[root@overcloud-controller-0 heat-admin]# cp /tmp/primary_disk
/var/spool/cisco/vsm/primary_disk
cp: overwrite /var/spool/cisco/vsm/primary_disk? y
[root@overcloud-controller-0 heat-admin]# exit
[heat-admin@overcloud-controller-0 ~]$ exit
```

Step 3 Repeat the above steps for the secondary VSM disk, ensure that you copy the secondary_disk only to the controllers hosting the secondary VSM.

Step 4 Shut down the following controller nodes in the specified order:

1. Controller node not hosting a VSM
2. Controller node with vsm-s
3. Controller node with vsm-p

Step 5 Restart the following controller nodes in the specified order:

1. Controller node with vsm-p
2. Controller node with vsm-s
3. Controller node not hosting a VSM

Step 6 Verify the setup.

a) Check the Nova status. For example:

```
[stack@n1kv-ospd-test ~]$ nova list
+--------------------------------------+------------------------+--------+------------+-------------+-------------------+
| ID | Name | Status | Task State | Power State | Networks |
+--------------------------------------+------------------------+--------+------------+-------------+-------------------+
| 0206d790-6b56-4b3e-a589-51e9150c39b1 | overcloud-compute-0 | ACTIVE | - | Running |
ctlplane=192.0.2.9 |
| c0140cc1-9f5f-4f34-a124-8a78f31586db | overcloud-controller-0 | ACTIVE | - | Running
| ctlplane=192.0.2.10 |
| ff0e325e-bfa7-4af2-836b-3acd12f72670 | overcloud-controller-1 | ACTIVE | - | Running
| ctlplane=192.0.2.11 |
| 364c507b-3c79-407e-86bd-f3d1eea5aa1f | overcloud-controller-2 | ACTIVE | - | Running
```

```
| ctlplane=192.0.2.12 |
+-------------------------+-------------------+-------+---------+---------+-----------------+
```

b) Check the Pacemaker status. For example:

```
[root@overcloud-controller-2 heat-admin]# pcs status | grep vsm
vsm-p (ocf::heartbeat:VirtualDomain): Started overcloud-controller-1
vsm-s (ocf::heartbeat:VirtualDomain): Started overcloud-controller-2
```

c) Check the VSM status. For example:

```
[stack@n1kv-ospd-test ~]$ ssh admin@192.0.2.50
Nexus 1000v Switch
Password:
vsm-p# sh mod
Mod Ports Module-Type Model Status
--- ----- -------------------------------- ------------------ ------------
1 0 Virtual Supervisor Module Nexus1000V active *
2 0 Virtual Supervisor Module Nexus1000V ha-standby
3 1022 Virtual Ethernet Module NA ok
4 1022 Virtual Ethernet Module NA ok
5 1022 Virtual Ethernet Module NA ok
6 1022 Virtual Ethernet Module NA ok
...
```

Basic Troubleshooting

Recovering from Multiple Controller Node Failure

The current HA model supports recovery from a single node failure. If there is multiple controller node failure, some VSM configuration may be lost. In such a scenario, to recover, use the backup of the VSM configuration saved at a remote location (See Overcloud Installation).

Procedure

Step 1 Log in to the primary VSM and verify both the active and standby VSMs using the **show module** command. If either primary or secondary VSM is not available, follow the procedure defined in Recovering VSM Failure before proceeding to the next step.

Step 2 Compare the running configuration with the configuration defined in the backup file. Use **show running-config** command to view the running configuration. If there are any discrepancies between the running configuration and backup configuration, run the missing configuration commands on the VSM.

Recovering from VSM Failure

Complete these steps to recover from a VSM failure caused due to multiple controller node failure. You will bring up a fresh VSM setup on new disks using the VSM configuration backup taken earlier.

Note The following procedure results in loss of any previous VSM configuration. If an up to date remote VSM Configuration backup is not available, please contact customer support for alternative recovery methods.

Procedure

Step 1 Disable both the primary and secondary VSM resources in pacemaker. For example:

```
[root@overcloud-controller-0 heat-admin]# pcs resource disable vsm-p
[root@overcloud-controller-0 heat-admin]# pcs resource disable vsm-s
```

Step 2 Log into nodes with active VSMs and shutdown the VMs. For example:

```
[root@overcloud-controller-0 heat-admin]# pcs status | grep vsm
vsm-p (ocf::heartbeat:VirtualDomain): Started overcloud-controller-2
vsm-s (ocf::heartbeat:VirtualDomain): Started overcloud-controller-0
[root@overcloud-controller-0 heat-admin]# virsh destroy vsm-s
[root@overcloud-controller-2 heat-admin]# virsh destroy vsm-p
```

Step 3 Log in to all the three controllers and run the following commands:

```
[root@overcloud-controller-0 heat-admin]# qemu-img create /var/spool/cisco/vsm/primary_disk
 4G
[root@overcloud-controller-0 heat-admin]# qemu-img create /var/spool/cisco/vsm/secondary_disk
 4G

[root@overcloud-controller-1 heat-admin]# qemu-img create /var/spool/cisco/vsm/primary_disk
 4G
[root@overcloud-controller-1 heat-admin]# qemu-img create /var/spool/cisco/vsm/secondary_disk
 4G

[root@overcloud-controller-2 heat-admin]# qemu-img create /var/spool/cisco/vsm/primary_disk
 4G
[root@overcloud-controller-2 heat-admin]# qemu-img create /var/spool/cisco/vsm/secondary_disk
 4G
```

Step 4 Enable both the primary and secondary VSMs in pacemaker. For example:

```
[root@overcloud-controller-0 heat-admin]# pcs resource enable vsm-p
[root@overcloud-controller-0 heat-admin]# pcs resource enable vsm-s
```

Step 5 Log in to the primary VSM and verify both the active and standby VSMs using the **show module** command.

Step 6 Compare the running configuration with the configuration defined in the backup file. Use **show running-config** command to view the running configuration. If there are any discrepancies between the running configuration and backup configuration, run the missing configuration commands on the VSM.