



Troubleshooting the Cisco Nexus 1000V Installation

This chapter describes how to identify and resolve problems related to installing the Cisco Nexus 1000V Switch for KVM on the Red Hat Enterprise Linux OpenStack Platform 7 (RHEL-OSP7).

Information About Cisco Nexus 1000V for KVM on the RHEL-OSP

The Cisco Nexus 1000V for KVM on the RHEL-OSP consists of these main components:

- **Virtual Ethernet Module (VEM)**—A software component that is deployed on each KVM host. Each VM on the host is connected to the VEM through virtual Ethernet (vEth) ports. The VEM is a hypervisor-resident component and is tightly integrated with the KVM architecture.
- **Virtual Supervisor Module (VSM)**—The management component that controls multiple VEMs and helps in the definition of VM-focused network policies. It is deployed either as a virtual appliance on any KVM host or on the Cisco Cloud Services Platform appliance. The VSM is integrated with OpenStack using the OpenStack Neutron plug-in.



Note This guide does not cover Cisco Nexus 1000V switch installation on the Cloud Services Platform.

- **RHEL-OSP**—Red Hat Enterprise Linux operating system with the Red Hat implementation of OpenStack Kilo. RHEL-OSP consists of services to control and manage computing, storage, and networking resources. These services provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud.

The Cisco Nexus 1000V for KVM uses Red Hat’s deployment management tool called Red Hat Enterprise Linux OpenStack Platform Director (also known as RHEL-OSPD) to install the Cisco Nexus 1000V for KVM on RHEL in an OpenStack cloud environment. The RHEL-OSP Director is based on the OpenStack-on-OpenStack (TripleO) project. The RHEL-OSP Director consists of two main components:

- **Undercloud**: The main director node that contains components for configuring and managing the OpenStack nodes that comprise the OpenStack environment (Overcloud). The main components of Undercloud provide functionality for environment planning, bare metal system control, and orchestration for OpenStack environment. For more information on Undercloud, see [Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage](#).

- Overcloud: The RHEL-OSP environment that is created using the Undercloud. The Overcloud comprises three main node types: controller nodes, compute nodes, and storage nodes. For more information on Overcloud, see [Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage](#).

Problems with Cisco Nexus 1000V Installation on OSP7

The following are symptoms, possible causes, and solutions for installation problems.

Symptom	Possible Causes	Solution
Overcloud fails to deploy due to a puppet failure.	The VSM management interface (N1000vVSMHostMgmtIntf) is set to the same interface as the provisioning interface.	Reconfigure the VSM to use a different management interface or refer to the Red Hat documentation and configure the provisioning interface on a bridge. Use the bridge for N1000vVSMHostMgmtIntf and set the N1000vExistingBridge parameter to true. These parameters are defined in the <code>cisco-n1kv-config.yaml</code> configuration file available at <code>/usr/share/openstack-tripleo-heat-templates/environments</code> .
Some VEMs cannot communicate with the VSM.	<ul style="list-style-type: none"> • Different vendor NICs are attached to your physical server. The management interface might be mapped to a different Ethernet interface other than the eth0 interface. • The parameters N1000vVSMHostMgmtIntf, N1000vVEMHostMgmtIntf, and N1000vExistingBridge are configured incorrectly in the configuration file, <code>/usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml</code>. • The parameter N1000vVEMHostMgmtIntf is configured incorrectly in the environment file. 	<p>Edit the parameter values to match the values in the configuration file (YAML file), or the environment file.</p> <p>If your compute and controller nodes have heterogeneous NIC ordering, you can leverage the custom configuration provided by the NodeDataLookup parameter to specify the different configurations. We recommend that you use the NodeDataLookup parameter to specify the configuration for controller nodes and use the N1000vVEMHostMgmtIntf parameter for the compute nodes.</p> <p>The number of controller nodes ranges from 1 to 3, whereas the number of compute nodes can expand over the life of your deployment. For details about how to leverage NodeDataLookup, see the <i>Cisco Nexus 1000V for KVM Installation Guide for Red Hat Enterprise Linux OpenStack Platform 7</i>.</p>

Symptom	Possible Causes	Solution
VSM bringup fails if the management (provisioning) interface of the controller node is not eth0.	<ul style="list-style-type: none"> Different vendor NICs are attached to your physical server. The management interface might be mapped to a different Ethernet interface other than the eth0 interface. The parameters N1000vVSMHostMgmtIntf and N1000vExistingBridge are configured incorrectly in the configuration file, <i>/usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml</i>. 	<ol style="list-style-type: none"> Configure the parameters in the configuration file <i>/usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml</i>. If the management interface is not eth0, configure the management interface of VSM on a separate uplink interface. Set the value of the N1000vVSMHostMgmtIntf parameter to an uplink interface name other than the name used for the controller node management interface. Also, set the N1000vExistingBridge parameter to <i>false</i>.
VSM boots to the loader prompt after multiple controllers nodes reboot ungracefully.	Multiple controller nodes fail simultaneously.	<p>Ensure that you have a backup of the latest VSM configuration at a remote location.</p> <ol style="list-style-type: none"> Disable pacemaker resources, such as the primary VSM (vsm-p) and the secondary VSM (vsm-s). Log in to nodes with active VSMS and shut down the VSM VMs. Log in to all three controllers and format the primary_disk and secondary_disk by using the qemu-img create disk-name 4G command at <i>/var/spool/cisco/vsm/</i>. Enable both the primary and secondary VSMS in pacemaker. <code>#pcs resource enable resource_id</code> Recover the missing VSM configuration from backup. <p>Note The VSM configuration might be lost during the recovery.</p>

Symptom	Possible Causes	Solution
VSMs go into a split-brain condition where the primary and secondary VSMs are in active - active state.	<ul style="list-style-type: none"> Layer 2 connectivity between the primary and secondary VSMs is lost. Multiple controller nodes fail simultaneously. 	<p>Ensure that you have a backup of the latest VSM configuration at a remote location.</p> <ol style="list-style-type: none"> Identify the primary and secondary VSM controller hosts by using the pcs status command. For example: <pre>[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</pre> Disable pacemaker resources, such as the primary VSM (vsm-p) and the secondary VSM (vsm-s). For example: <pre>#pcs resource disable resource_id</pre> Log in to nodes with active VSMs and shut down the VSM VMs. Log in to all three controllers and format the primary_disk and secondary_disk by using the qemu-img create disk-name 4G command at <i>/var/spool/cisco/vsm/</i>. Enable both the primary and secondary VSMs in pacemaker. <pre>#pcs resource enable resource_id</pre> Recover the missing VSM configuration from backup. <p>Note If you cannot log in to one of the VSMs via virsh console, use the peer mac-addresses clear command on the active VSM accessible through the virsh console command.</p> <p>Note The VSM configuration might be lost during the VSM split-brain recovery.</p>

Symptom	Possible Causes	Solution
VSM standby is not running.	Pacemaker cannot reinitiate a primary standby or secondary standby node.	<ol style="list-style-type: none"> 1. Check whether the primary or secondary node is in standby mode. Log in to the active VSM and run the show redundancy status command. 2. Run the pcs resource cleanup [<i>vsm-pl/vsm-s</i>] command from one of the controller nodes for the standby VSM. 3. Check the pacemaker status: <pre>[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</pre>
VEM configuration is missing.	<ul style="list-style-type: none"> • Per node specified in the NodeDataLookup configuration is not correctly applied to the node. • The Puppet apply command to apply configurations failed to run. 	<ol style="list-style-type: none"> 1. Log in to the node with the VEM configuration problem. 2. Run the dmidecode --s system-uuid command to retrieve the System UUID. 3. Open the <System-UUID>.json file and confirm whether all configuration parameters expected for the configuration file, <i>n1kv.conf</i>, are present. 4. If the parameter values are incorrect, go to the Undercloud and reconfirm the VEM override parameter NodeDataLookup in the configuration file <i>/usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml</i>. 5. Perform a heat stack update by redeploying Overcloud to have the latest configurations on the respective nodes.
Virtual Ethernet interfaces on VSM corresponding to the router ports on VEM flap continuously.	The Neutron l3_ha parameter is set to <i>True</i> and one or more OpenStack controller nodes (in HA mode) are down or are in inconsistent state.	<ol style="list-style-type: none"> 1. Edit the configuration file, <i>./etc/neutron/neutron.conf</i>, on the OpenStack controller node and set the l3_ha parameter to <i>false</i> and the allow_automatic_l3agent_failover parameter to <i>true</i>. 2. Restart the neutron-server service and the neutron-l3-agent service. 3. Repeat Step 1 and Step 2 on all OpenStack controller nodes in HA. 4. Clean up the router ports by manually deleting and recreating the router ports.

Symptom	Possible Causes	Solution
Inter-VLAN traffic has stopped for VMs.	The Neutron <code>allow_automatic_l3agent_failover</code> parameter is set to <i>False</i> and one or more OpenStack controller nodes (in HA mode) are down or are in inconsistent state, causing an automatic migration of <code>l3_agent</code> ports failure.	<ol style="list-style-type: none"> Edit the configuration file, <code>./etc/neutron/neutron.conf</code>, on the OpenStack controller node and set the <code>allow_automatic_l3agent_failover</code> parameter to <i>true</i>. Restart the <code>neutron-server</code> service and the <code>neutron-l3-agent</code> service. Repeat Step 1 and Step 2 on all OpenStack controller nodes in HA.
Uplink Ethernet and Virtual Ethernet (VTEP) ports show NoPortProfile state on VSM after deploying using OSP7.	<ul style="list-style-type: none"> Port profiles required to bring up ports on VSM are not defined on the VSM. The port profile names defined in the heat template (<code>/usr/share/openstack-tripleo-heat-templates/environments/cisco-n1kv-config.yaml</code>) do not match the port profile names defined on the VSM. 	Create Ethernet and VTEP Virtual Ethernet port profiles with the same name as defined in the heat template (<code>cisco-n1kv-config.yaml</code>) on the Undercloud node.
VSM shows the VEM module as offline or a VEM module is missing from VSM <code>show module</code> command output after the deployment through OSP7.	The Cisco Nexus 1000V VEM service is not running on the respective VEM module.	<p>Verify the state of the Cisco Nexus 1000V VEM service on the VEM host by using the service nexu1000v status command.</p> <p>If the service is not running, restart it by using the service nexu1000v start command.</p>
	The Cisco Nexus 1000V VEM cannot communicate with the VSM due to an incorrect networking configuration.	Revisit the network planning. You can leverage the <code>NodeDataLookup</code> parameter to provide node-specific configuration for a single node or class of nodes to enable heterogeneous deployment.
An error is observed on the Openstack controller node when you try to access the policy profiles pushed via the Cisco Nexus 1000V VSM using the neutron cisco-policy-profile-list command.	The configuration file, <code>./etc/neutron/Neutron.conf</code> , is missing a value for the <code>service_plugins</code> parameter.	<p>Update the <code>neutron.conf</code> configuration file with the <code>service_plugins</code> parameter value:</p> <pre>service_plugins =router,cisco_n1kv_profile</pre> <p>After updating the parameter value on all controllers, restart the <code>neutron-server</code> on all controllers in HA mode.</p>