Overview

This chapter contains the following sections:

Information About Cisco Nexus 1000V for KVM on Red Hat Enterprise Linux OpenStack Platform

The Cisco Nexus 1000V for KVM on Red Hat Enterprise Linux OpenStack Platform (RHEL-OSP) consists of these main components:

- Virtual Ethernet Module (VEM)—A software component that is deployed on each KVM host. Each virtual machine (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 Plugin.

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

- RHEL-OSP—Red Hat Enterprise's Linux operating system with the Red Hat's implementation of the 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.

Information About the Red Hat Enterprise Linux OpenStack Platform Director

The Cisco Nexus 1000V for KVM uses Redhat'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 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 of 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.

Red Hat Enterprise Linux OpenStack Platform provides the foundation to build a private or public Infrastructure-as-a-Service (IaaS) cloud on top of Red Hat Enterprise Linux. It supports High Availability (HA) in an OpenStack Platform environment. The HA support is implemented using a Controller node cluster and a Pacemaker (cluster resource manager). For more information on HA support, see Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage.

Supported Network Topologies

Topology with OpenStack in Standalone Mode

The Cisco Nexus 1000V for KVM can be deployed with OpenStack in standalone mode (not functioning in high-availability (HA) mode). However, Cisco recommends that you always deploy the VSM in active/standby HA mode. The following topology diagram shows standalone mode deployment topology.

Figure 1: Topology with OpenStack in Standalone Mode

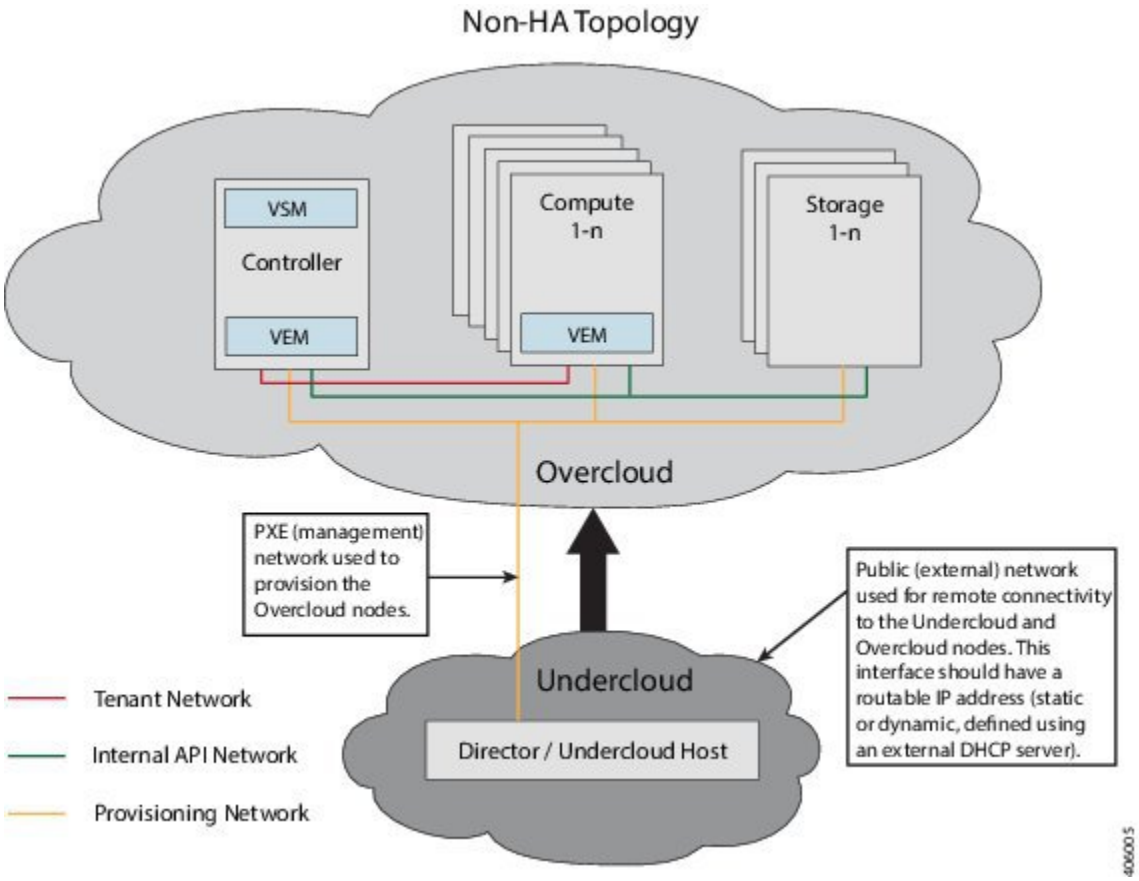

Topology with OpenStack in High-Availability Mode

The Cisco Nexus 1000V for KVM can be deployed with OpenStack in high-availability (HA) mode. We recommend that you always deploy the VSM in active/standby HA mode. The following topology diagram shows OpenStack in HA mode deployment.

Figure 2: Topology with OpenStack in High-Availability Mode

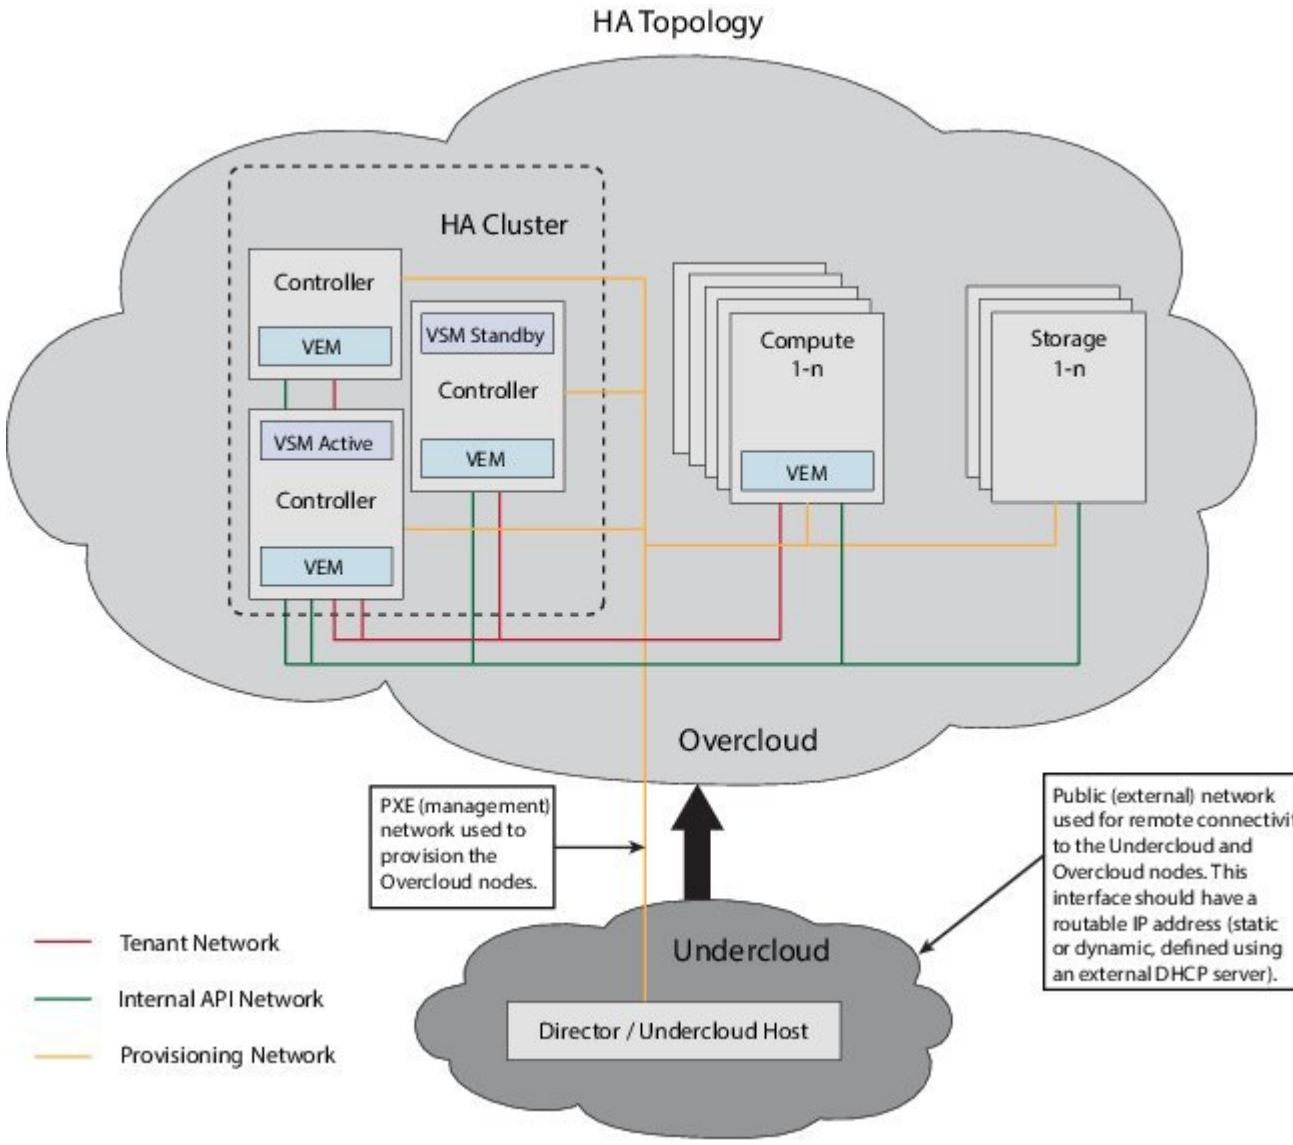

OpenStack Hosts and Services

Hosts and Services Used in OpenStack Standalone Deployments

The following table lists and describes the primary OpenStack hosts and services for Nexus1000V deployment that are needed when deploying OpenStack in standalone mode:

Hosts	OpenStack Service
Director or Undercloud Host	Deploys the OpenStack services
Controller	Neutron Server
	Database (MySQL)
	Messaging (RabbitMQ/QPID)
	Heat
	Ceilometer
	Keystone
	Glance
	Nova
	Cinder
	Horizon or OpenStack Dashboard
	Neutron Layer 3 Agent
	Neutron DHCP Agent
	Neutron Metadata Agent
	Cisco Nexus 1000V Virtual Ethernet Module (VEM)
Compute	Ceilometer Agent
	Cisco Nexus 1000V Virtual Ethernet Module (VEM)
	Nova-compute

Hosts and Services Used in OpenStack High-Availability Deployments

The primary OpenStack hosts and services that you need when deploying OpenStack in HA mode are same as those in a standalone mode. For more information, see Hosts and Services Used in OpenStack Standalone Deployments, on page 5.

OpenStack HA deployment is defined by the number of controllers. A single controller indicates that it is a non-HA setup, while three controllers indicate that it is an OpenStack HA setup.

Note For pacemaker setup, you need more than half of configured or expected controllers to be active state. Hence, in a three controller cluster, you need to have at least two controllers in active state. For more information about pacemaker, see: Pacemaker Cluster Documentation.

Heat Templates

The orchestration required to bring up Cisco Nexus 1000V Switch instance in the OpenStack environment is implemented using Heat templates. Heat is one of the main components of the OpenStack orchestration program. The Heat template is used to describe the deployment of complex cloud applications on the OpenStack platform. The Heat Orchestration Template (HOT) is written as a structured YAML text file and defines the OpenStack resources such as compute resources and network resources. The Heat template also defines the configuration information for each OpenStack resource. The Heat engine parses and implements the Heat template. A Heat template contains three main sections: parameters, resources, and output.

The following table describes Heat parameters, VSM puppet parameters, and plugin parameters.

Note These parameters are defined in the default configuration file named cisco-n1kv.yaml, available at `/usr/share/openstack-tripleo-heat-templates/puppet/extraconfig/pre_deploy/controlle`
You can override the default parameter values in the cisco-n1kv.yaml file. To override the default values, specify the new values in the configuration file named cisco-n1kv-config.yaml, available at `/usr/share/openstack-tripleo-heat-templates/environments`.

Table 1: Heat Parameters, VSM Puppet Parameters, and Plugin Parameters

Parameter	Type	Default Value	Description
N1000vVSMIP	string	192.0.2.50	IP address of the Virtual Supervisor Module (VSM).
N1000vVSMDomainID	number	100	Domain ID of the VSM.
N1000vVEMHostMgmtIntf	string	br-ex	Name of the existing bridge for the VSM or the name of the physical interface to be configured on the VSM bridge.

Parameter	Type	Default Value	Description
N1000vUplinkProfile	string	{eth1: system-uplink,}	Uplink interfaces managed by the VEM. You must also specify the uplink port profile that configures these interfaces. For example, "{eth1: system-uplink,}".
N1000vVtepConfig	string	undefined	Virtual tunnel interface configuration for the VXLAN tunnel endpoints. For example, {"vtpe1": {"profile": "virtprof", "ipmode": "dhcp"},}
N1000vVEMFastpathFlood	string	enable	Enable or disable the broadcasting and flooding of unknown packets in the kernel module.
N1000vVSMVersion	string	5.2.1.SK3.2.2b-1	Version of VSM image that is used for the deployment.
NeutronServicePlugins	string		Value for Neutron service plugins. This should be set to router, cisco_n1kv_profile to enable the policy profile service plugin along with the default router service plugin.
NeutronTypeDrivers	string		Neutron type drivers to be enabled. The value should include either VLAN or VXLAN or both, based on the network required.
NeutronCorePlugin	string		Neutron core plugin value. This should be set to neutron.plugins.ml2.plugin.Ml2Plugin to enable the ML2 core plugin.
NeutronMechanismDrivers	string		Mechanism drivers to be loaded for the Neutron ML2 plugin. To enable Cisco Nexus 1000V, include cisco_n1kv in the parameter list, else Cisco Nexus 1000V configuration will not be loaded.
NeutronDhcpAgentsPerNetwork	number	1	Specify the number of DHCP agents configured for host a network.
NeutronAllowL3AgentFailover	boolean	true	Specify whether to allow automatic rescheduling of routers from dead L3 agents with admin_state_up (set to True) to alive agents.
NeutronL3HA	boolean	false	Specify whether high availability for the virtual routers is supported or not.
NodeDataLookup	string		Define host specific overrides. For detailed information about this parameter, see Extracting System UUID to Configure Nodes.

Parameter	Type	Default Value	Description
N1000vVSMHostMgmtIntf	string		Interface or bridge on the controller node that the VSM uses for connectivity. If it is a bridge, the parameter, N1000vExistingBridge, should be set to true and patch ports are created between the bridge on which the VSM ports are created and the existing bridge. If it is an interface, the bridge on which the VSM ports are created use the interface as an uplink port for the VSM ports. **Note** You cannot be set this parameter for the provisioning interface because the VSM installation may interrupt connectivity.
N1000vExistingBridge	boolean	true	Specify whether the VSM is brought up using an existing bridge or using a new bridge. If this parameter is set to true, the VSM is brought up on an existing bridge, specified in the N1000vVSMHostMgmtIntf parameter. If set to false, the VSM is brought up on new bridge that is created using the interface defined in the N1000vVSMHostMgmtIntf parameter.
N1000vVSMPassword	string	Password	Password for the Cisco Nexus 1000V VSM.
N1000vMgmtNetmask	string	255.255.255.0	Netmask for the Cisco Nexus 1000V VSM.
N1000vMgmtGatewayIP	string	127.0.0.2	Gateway for the Cisco Nexus 1000V VSM.
N1000vPacemakerControl	boolean	true	Specify whether the VSM is managed by the pacemaker controller or not. If this parameter is set to true, the VSM is managed by pacemaker controller. **Note** Set this parameter to false for non-HA setup.
N1000vVSMUser	string	admin	Username for all the configured Cisco Nexus1000V VSMs.
N1000vPollDuration	number	60	Cisco Nexus 1000V policy profile polling duration in seconds.
N1000vHttpPoolSize	number	5	Number of threads to use to make HTTP requests.
N1000vHttpTimeout	number	15	HTTP timeout, in seconds, for connection to the Cisco Nexus1000V VSMs.

Parameter	Type	Default Value	Description
N1000vSyncInterval	number	300	Time interval between consecutive neutron-VSM synchronizations.
N1000vMaxVSMRetries	number	2	Maximum number of retry attempts for VSM REST API.

Note For detailed information about Heat templates, see the Red Hat Enterprise Linux OpenStack Platform 7 Director Installation and Usage