



## Overview

---

This chapter contains the following sections:

- [Information About the Installation, page 1](#)
- [Juju Charms, page 2](#)
- [Configuration Files, page 3](#)
- [Types of Nodes, page 4](#)
- [Supported Nodes, page 6](#)
- [Network Connectivity Between OpenStack Services, page 6](#)
- [Supported Topologies, page 7](#)

## Information About the Installation


Cisco Nexus 1000V for KVM uses the Ubuntu Metal as a Service (MAAS) and Juju tools to facilitate the installation of OpenStack and Cisco Nexus 1000V for KVM. A description of each is as follows:

- **Metal as a Service (MAAS)**—Tool that sets up and manages the physical infrastructure on which services are deployed.
- **Juju**— Tool that deploys services, such as OpenStack and the Cisco Nexus 1000V for KVM services to your physical or virtual environment. Juju provides the installation logic (Juju charm) and software packages (Debian packages) to deploy the Cisco Nexus 1000V for KVM.
- **OpenStack**—Scalable cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter.
- **Cisco Nexus 1000V for KVM**— Distributed virtual switch (DVS) that works with several different hypervisors. This DVS version is integrated with the Ubuntu Linux Kernel-based virtual machine (KVM) open source hypervisor.

You need to deploy MAAS, Juju, and OpenStack before you can deploy the Cisco Nexus 1000V for KVM.

The following figure shows the installation process:

**Figure 1: Installation Process Overview**



## Juju Charms

The charm logic to install the Cisco Nexus 1000V-specific OpenStack Networking plugins and the Cisco Nexus 1000V for KVM services are packaged together in a charm Debian package called `jujucharm-n1k_5.2.1.sk1.2.2.YYYYMMDDhhmm-1_amd64.deb`. The MAAS node downloads these charms to the cluster nodes to facilitate the installation of their respective services.

- OpenStack charms
  - nova-cloud-controller
  - quantum-gateway
  - openstack-dashboard
  - nova-compute
  - ceph
  - ceph-radosgw
  - cinder
  - glance
  - keystone
  - mysql
  - rabbitmq-server
  - swift-proxy
  - swift-storage
- Cisco Nexus 1000V for KVM charms
  - vsm
  - vem
  - vxlan-gateway

Juju deploys each service using its respective charm and two configuration files: its own `config.yaml` file (comes with the charm and serves as its parameter data modeling and default value settings) and a global configuration file (a file that you create for defining any deployment-specific parameters).



**Note** Treat the `config.yaml` file as a read-only file. If you want to add any deployment-specific parameter changes, use the global configuration file.

For more information about the global configuration file file, see [Preparing for Installation](#).

The VSM charm helps to deploy the Virtual Supervisor Module (VSM) as a virtual machine on an Ubuntu KVM server node in a MAAS environment. This charm does not install the VSM on a Cisco Nexus Cloud Service Platform. On a Cloud Service Platform, VSM is installed through the standard method without the Juju charm. For more information, see [Installing VSM on the Cisco Nexus Cloud Services Platform](#).

The VEM charm (`vem`) is a subordinate charm under two primary charms: `nova-compute` and `quantum-gateway`. You deploy the `vem` first; however, the real instantiation of VEM happens when you add relationships between the `vem` charm and either of the two primary charms.

When you deploy the Virtual Ethernet Module (VEM), you must add relationships between the VEM and the Nova-compute node and between VEM and the quantum-gateway so that whenever these nodes are deployed, the VEM service is also deployed automatically. When this relationship is built, the VEM charm installs a VEM on every nova-compute and quantum-gateway nodes.

The VEM can be configured by using the `config.yaml` file and the global configuration file that you provide. In addition, you can define specific information for a host by using a variable named "mapping" which takes in the content of a server node mapping file (in YAML format and named, for example, `mapping.yaml`). This file is delineated by the host ID (for example, `maas-node-3`) where you specify each server node variable value that is different from the other server nodes. The most common usage of this mapping file is for configuring the VXLAN Tunnel Endpoints (VTEPs) to implement the Virtual Extensible Local Area Network (VXLAN) feature. VTEPs are host specific and require that you define host-specific values.

## Configuration Files

The `config.yaml` file defines the whole set of configurable parameters for each charm. It defines each parameter's type, default value, and corresponding description. You do not modify this file. For an example of this file, see [Sample Global Configuration File](#).

To define values for parameters that are specific to your deployment, you need to create a global configuration file. You separate the charms into sections and provide the corresponding parameters and values in each section. For more information, see [Preparing the Configuration and Mapping Files](#).

For the Cisco Nexus 1000V-related OpenStack charms listed below, you need to modify the global configuration file within each individual charm's section with this provision: `openstack-origin:deb https://user:password@private-ppa.launchpad.net/springfield-team/havana-staging-2013.2.2/ubuntu precise main|key-id`.

- `nova-cloud-controller`
- `nova-compute`
- `openstack-dashboard`
- `quantum-gateway`

**Note**

To deploy the VEMs with their specific parameters, you need to create a host-mapping file, which is deployed with the other configuration files when you deploy the VEM service. For more information see, [Cisco Nexus 1000V for KVM VEM Charm Parameters](#).

## Types of Nodes

A node in MAAS architecture can be a physical server or a virtual machine. MAAS does not differentiate between these two.

The Cisco Nexus 1000V for KVM deployment has the following five node types:

- Infrastructure nodes (MAAS server node and Juju node)
- OpenStack service nodes
- Network nodes
- Nova compute nodes
- Virtual Supervisor Module (VSM) nodes

### MAAS Server Node

The MAAS server node is the cluster controller where the cluster provisioning and service management are performed. It provides the following functions:

- DHCP and Preboot Execution Environment (PXE) boot services
- DNS and DNS forwarding service
- Cluster controller service for enlisting and delisting the cluster nodes
- Debian package proxy
- Juju client when the Juju command is issued

### Juju Node

The Juju node is the server node where the following functions are provided:

- Serves the Juju request from clients (for example, from the MAAS server)
- Service scheduling
- Service staging
- Establishing and maintaining relationships between services

### OpenStack Service Nodes

MAAS OpenStack services can be deployed as VMs on a single server or on different servers. In either case, they are referred to as OpenStack service nodes. MAAS OpenStack provides the following services through these service charms:

Services	Charms
Cloud controller and Cisco Nexus 1000V VXLAN Gateway Glance image retrieval and repository (Nexus 1000V)	nova-cloud-controller vxlan-gateway (subordinate charm to nova-cloud-controller)
Identity service	keystone
Unified, distributed storage service	ceph
Amazon Simple Storage Service (S3), Swift-compatible HTTP gateway for online object storage on top of a Ceph cluster (RADOS Gateway)	ceph-radosgw
Volume service	cinder
Image service	glance
Django-based web user-interface for administering the OpenStack nodes	openstack-dashboard
Proxy for object storage service	swift-proxy
Object storage service	swift-storage
SQL database service	mysql
Advanced Message Queuing Protocol (AMQP) messaging service	rabbitmq

### Network Nodes

The network nodes host the OpenStack networking services for your Cisco Nexus 1000V for KVM deployment, including all of the OpenStack service nodes. The network nodes provide the following services through these service charms:

Services	Charms
Neutron DHCP and Layer 3 agents	quantum-gateway vem (subordinate charm to quantum-gateway)

### Nova Compute Nodes

The Nova compute nodes host your virtual machines (VMs), including any VXLAN Gateway VMs that you have deployed as VMs. The Nova compute nodes provide the following service through these service charms:

Service	Charms
Nova compute service	nova-compute vem (subordinate charm to nova-compute)

### VSM Nodes

The VSM can be hosted on a dedicated server node or on a Cisco Nexus 1110 as a Virtual Service Blade (VSB). The VSM node provides the following service through this service charm:

Service	Charm
Distributed virtual switch management and control of multiple Virtual Ethernet Modules (VEMs).	vsm

## Supported Nodes

The Cisco Nexus 1000V for KVM supports the following nodes in a MAAS OpenStack cluster:

- One MAAS server node
- One Juju node
- One or more OpenStack service nodes
- One or more Nova-compute nodes
- Zero, one, or more network nodes
- One or two VSM nodes if you deploy the VSM as a VM. Zero VSM nodes when you deploy the VSM as a VSB in a Cisco Cloud Services Platform.

## Network Connectivity Between OpenStack Services

A MAAS deployment requires the following four functional networks:

- Management network—Provides internal communication between the OpenStack components. All IP addresses on this network need to be reachable only within the data center.
- Data network—Provides communication between VMs in the cloud. The IP addressing requirements of this network depend on the OpenStack network plugin that your deployment uses.
- Public network—Provides a node with public internet access. If required, IP addresses on this network need to be reachable from the internet.
- Intelligent Platform Management Interface (IPMI) network—Manages the power sequence for all of the nodes in the cluster.

System administrators retain the option to collapse network boundaries based on the physical setup of their datacenter. For example, you can fold the management network into the data network.

# Supported Topologies

Cisco Nexus 1000V for KVM supports both physical and virtual machine topologies.

## Physical Deployment

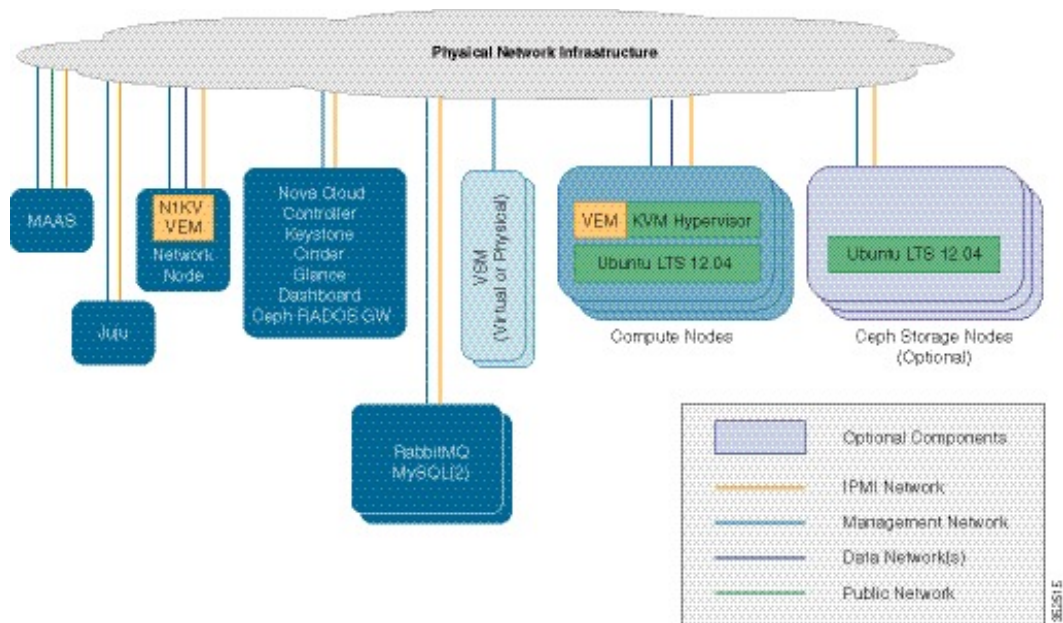
In this deployment model, the services are deployed on a baremetal server rather than as virtual machines. (See the following figure.) The primary benefit of this deployment model is the performance of the OpenStack services as well as the network nodes. When a Layer 3 agent with high performance is required, this deployment model is recommended.



**Note**

In all deployment models, we recommend that you deploy the VSM in high availability (HA) mode.

**Figure 2: Topology of Physical Deployment**



## Virtual Machine Deployment

The simplest deployment model is one where all of the OpenStack, MAAS, and Juju services are deployed on a single physical server as virtual machines (see the following figure). The network node (used for offering the DHCP/Layer 3 Agent service) is also deployed on this server as a virtual machine. Multiple network node instances can be created for scaling the agent ports. However, the capacity of the Layer 3 agents is limited in this solution because all the Layer 3 agents are running on the same physical server.

In this solution, a VSM can be deployed as a virtual machine or on a Cloud Services Platform. Similarly, the network node can be deployed as a VM or on a physical server. However, we recommend that you deploy the network node on a physical server, as shown in the figure.



**Note** In all deployment models, we recommend that you deploy the VSM in high availability (HA) mode.

**Figure 3: Topology of Virtual Machine Deployment**

