



## 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 and Juju before you can deploy OpenStack with the Cisco Nexus 1000V for KVM.

The following figure shows the installation process:

**Figure 1: Installation Process Overview**



If you want to deploy OpenStack with the Cisco Nexus 1000V for KVM on a smaller scale for testing purposes, see [Setup Script](#).

## 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_sk3.1.1YYYYMMDDhhmm-1_amd64`. This package contains the following OpenStack and Cisco Nexus 1000V for KVM charms:

- nova-cloud-controller
- quantum-gateway
- openstack-dashboard
- nova-compute
- vsm
- vem
- vxlan-gateway

The Metal-as-a-Service (MAAS) node downloads the OpenStack and Cisco Nexus 1000V for KVM charms to the cluster nodes to facilitate the installation of their respective services.

You must download each of the following required charms separately from the Charm store:

- ceph
- ceph-radosgw
- cinder
- glance
- keystone
- mysql
- rabbitmq-server
- swift-proxy

- swift-storage

Juju deploys each service using its respective charm and two configuration files: its own YAML Ain't Markup Language (YAML) configuration file (called `config.yaml` that 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, see [Preparing for Installation](#).

The Virtual Supervisor Module (VSM) charm helps to deploy the 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 Virtual Ethernet Module (VEM) charm is a subordinate charm under two primary charms: `nova-compute` and `quantum-gateway`. You deploy the VEM first; however, the real instantiation of the VEM happens when you add relationships with the corresponding primary charms.

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

Each charm's parameters are defined in a configuration file in YAML format.

**Note**

YAML (rhymes with “camel”) is a human-friendly, cross language, Unicode based data serialization language designed around the common native data structures of agile programming languages. It is broadly useful for programming needs ranging from configuration files to Internet messaging to object persistence to data auditing.

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 with this provision: `openstack-origin: ppa:cisco-n1kv/icehouse-updates`.

- 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 bootstrap node)
- OpenStack service nodes
- Network nodes (also known as Quantum-gateway 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 Bootstrap Node

The Juju bootstrap 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

If you decide to deploy any of the following OpenStack services, they can be deployed using their corresponding charms:

Services	Charms
Cloud controller	nova-cloud-controller
Loads Cisco Nexus 1000V VXLAN Gateway image to Glance	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
Percona XtraDB Cluster, which provides an active/active MySQL-compatible alternative that is implemented using the Galera synchronous replication extensions	percona-cluster

### 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 Metal-as-a-Service (MAAS) OpenStack cluster:

- One MAAS server node
- One Juju bootstrap node
- One or more OpenStack service nodes
- One or more Nova-compute nodes
- Zero, one, or more network nodes
- One or two Virtual Supervisor Module (VSM) nodes if you deploy the VSM as a virtual machine (VM). Zero VSM nodes when you deploy the VSM as a Virtual Service Blade (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.

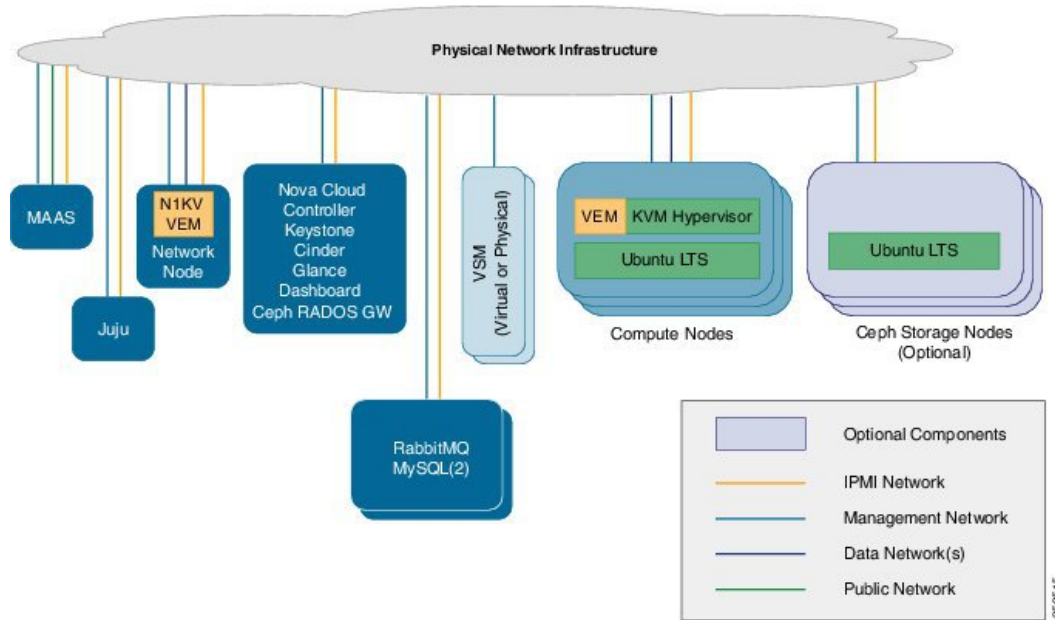
### Physical Deployment

In this deployment model, the services are deployed on bare metal servers rather than as virtual machines. See the figure below. The primary benefit of this deployment model is the performance improvement 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.

The VSM can be deployed in active/standby high availability (HA) mode using the Juju charms, or it can be manually brought up on a Cisco Nexus 1010 or 1110 Virtual Service Appliance. In all deployment models, we recommend that you deploy the Virtual Supervisor Module (VSM) in HA mode.

You can also deploy OpenStack in HA mode. For more information, see [OpenStack High Availability Deployment](#), on page 8.

**Figure 2: Topology of Physical Deployment**



## Virtual Machine Deployment

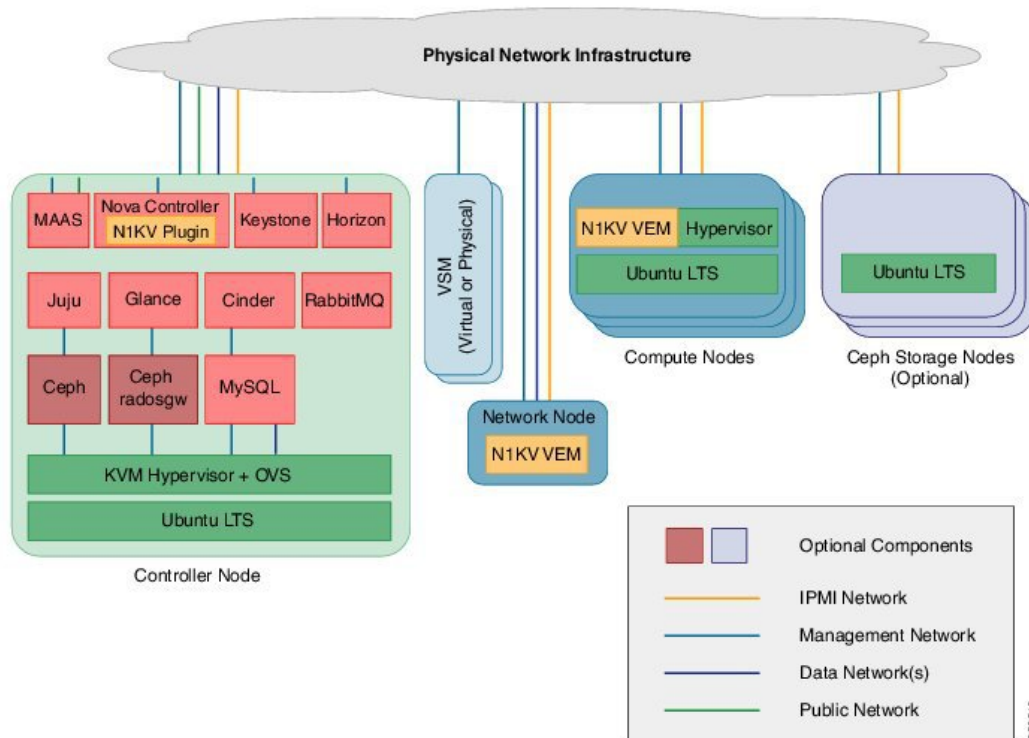
The simplest deployment model is one where all of the OpenStack, Metal-as-a-Service (MAAS), and Juju services are deployed on a single physical server as virtual machines (VMs). See the figure below. The network node (used for offering the DHCP/Layer 3 Agent service) is also deployed on this server as a VM. 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 Virtual Supervisor Module (VSM) can be deployed as a VM or on a Cloud Services Platform (CSP). 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.

The VSM can be deployed in Active/Standby HA mode using the Juju charms, or it can be manually brought up on a Cisco Nexus 1010 or 1110 Virtual Service Appliance. In all deployment models, we recommend that you deploy the VSM in high availability (HA) mode.

You can also deploy OpenStack in HA mode. For more information, see [OpenStack High Availability Deployment](#), on page 8.

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



## OpenStack High Availability Deployment

You can deploy the Ubuntu Openstack portion of Cisco Nexus 1000V in High Availability (HA) mode using Juju charms. No additional Cisco Nexus 1000V configuration is required. For information about the



requirements for deploying Ubuntu OpenStack in HA mode, see the documentation at the following location: <https://wiki.ubuntu.com/ServerTeam/OpenStackHA>.

The reason that the Cisco Nexus 1000V requires no additional configuration is that the Cisco Nexus Plug-in for OpenStack Neutron resides in the Nova-cloud-controller (NCC), which provides the API endpoints for the Nova and Neutron services. As APIs are stateless, the NCC nodes can be scaled horizontally by load-balancing the requests across all of the available nodes. Therefore, only one active NCC node processes the requests, and, when necessary, the node sends the requests to the Virtual Supervisor Module (VSM).

