



Overview

This chapter contains the following sections:

- [Information About Cisco Nexus 1000V, page 1](#)
- [Information About System Port Profiles and System VLANs, page 4](#)
- [Installation Overview, page 5](#)
- [Recommended Topologies, page 6](#)

Information About Cisco Nexus 1000V

The Cisco Nexus 1000V is a distributed virtual switch solution that is fully integrated within the VMware virtual infrastructure, including VMware vCenter, for the virtualization administrator. This solution offloads the configuration of the virtual switch and port groups to the network administrator to enforce a consistent data center network policy.

The Cisco Nexus 1000V is compatible with any upstream physical access layer switch that is compliant with the Ethernet standard, including the Catalyst 6500 series switch, Cisco Nexus switches, and switches from other network vendors. The Cisco Nexus 1000V is compatible with any server hardware that is listed in the [VMware Hardware Compatibility List \(HCL\)](#).



Note

We recommend that you monitor and install the patch files for the VMware ESXi host software.

The Cisco Nexus 1000V has the following components:

- Virtual Supervisor Module (VSM)—The control plane of the switch and a VM that runs Cisco NX-OS.
- Virtual Ethernet Module (VEM)—A virtual line card that is embedded in each VMware vSphere (ESXi) host. The VEM is partly inside the kernel of the hypervisor and partly in a user-world process, called the VEM Agent.

Virtual Supervisor Module

The VSM, along with the VEMs that it controls, performs the following functions for the Cisco Nexus 1000V system:

The VSM uses an external network fabric to communicate with the VEMs. The VSM runs the control plane protocols and configures the state of each VEM, but it never forwards packets. The physical NICs on the VEM server are the uplinks to the external fabric. VEMs switch traffic between the local virtual Ethernet ports that are connected to the VM vNICs but do not switch traffic to other VEMs. Instead, a source VEM switches packets to the uplinks that the external fabric delivers to the target VEM.


A single Cisco Nexus 1000V instance, including dual-redundant VSMs and managed VEMs, forms a switch domain. Each Cisco Nexus 1000V domain within a VMware vCenter Server must be distinguished by a unique integer called the domain identifier.

A single VSM can control up to 250 VEMs.

See the *Cisco Nexus 1000V Resource Availability Reference* for more information about scale limits.

The Cisco Nexus 1000V architecture is shown in the following figure.

Figure 1: Cisco Nexus 1000V Architecture



332119

Virtual Ethernet Module

Each hypervisor is embedded with one VEM that replaces the virtual switch by performing the following functions:

- Advanced networking and security

- Switching between directly attached VMs
- Uplinking to the rest of the network

**Note**

Only one version of the VEM can be installed on an ESXi host at any time.

**Note**

Cisco Nexus 1000V VEM does not support ESXi custom TCP/IP stack and control traffic through the custom TCP/IP stack.

In the Cisco Nexus 1000V, the traffic is switched between VMs locally at each VEM instance. Each VEM also interconnects the local VM with the rest of the network through the upstream access-layer network switch (blade, top-of-rack, end-of-row, and so forth). The VSM runs the control plane protocols and configures the state of each VEM accordingly, but it never forwards packets.

In the Cisco Nexus 1000V, the module slots are for the primary module 1 and secondary module 2. Either module can act as active or standby. The first server or host is automatically assigned to module 3. The network interface card (NIC) ports are 3/1 and 3/2 (vmmnic0 and vmmnic1 on the ESXi host). The ports to which the virtual NIC interfaces connect are virtual ports on the Cisco Nexus 1000V where they are assigned with a global number.

Information About VSM-to-VEM Communication

The VSM and the VEM can communicate over a Layer 2 network or a Layer 3 network. These configurations are referred to as Layer 2 or Layer 3 control modes.

Layer 3 Control Mode

Layer 3 control mode is the preferred method of communication between the VSM and the VEMs. In Layer 3 control mode, the VEMs can be in a different subnet than the VSM and from each other. Active and standby VSM control ports should be Layer 2 adjacent. These ports are used to communicate the HA protocol between the active and standby VSMs.

Each VEM needs a designated VMkernel NIC interface that is attached to the VEM that communicates with the VSM. This interface, which is called the Layer 3 Control vmknic, must have a system port profile applied to it (see [System Port Profiles, on page 4](#) and [System VLANs, on page 5](#)), so the VEM can enable it before contacting the VSM.

For a sample topology diagram, see [Topology for Layer 3 Control Mode, on page 6](#).


For more information about Layer 3 control mode, see the “Configuring the Domain” chapter in the *Cisco Nexus 1000V System Management Configuration Guide*.

Layer 2 Control Mode

In Layer 2 control mode, the VSM and VEMs are in the same subnet. You can install the VSM and VEMs on different ESXi hosts or on the same ESXi host. This figure shows a VSM and VEM that are running on the same host in Layer 2 control mode.

For a sample topology diagram showing Layer 2 control mode, see [Topology for Layer 2 Control Mode](#), on page 7.

Figure 2: VSM and VEM on the Same Host in Layer 2 Control Mode



Information About System Port Profiles and System VLANs

System Port Profiles

System port profiles can establish and protect ports and VLANs that need to be configured before the VEM contacts the VSM.

When a server administrator adds a host to a DVS, its VEM must be able to contact the VSM. Because the ports and VLANs used for this communication are not yet in place, the VSM sends a minimal configuration, including the system port profiles and system VLANs, to vCenter Server, which then propagates it to the VEM.

When configuring a system port profile, you assign VLANs and designate them as system VLANs. The port profile becomes a system port profile and is included in the Cisco Nexus 1000V opaque data. Interfaces that use the system port profile, which are members of one of the defined system VLANs, are automatically enabled and forward traffic when the VMware ESX starts even if the VEM does not have communication with the VSM. The critical host functions are enabled even if the VMware ESX host starts and cannot communicate with the VSM.

**Caution**

VMkernel connectivity can be lost if you do not configure the relevant VLANs as system VLANs.

System VLANs

You must define a system VLAN in both the Ethernet and vEthernet port profiles to automatically enable a specific virtual interface to forward traffic outside the ESX host. If the system VLAN is configured only on the port profile for the virtual interface, the traffic is not forwarded outside the host. Conversely, if the system VLAN is configured only on the Ethernet port profile, the VMware VMkernel interface that needs that VLAN is not enabled by default and does not forward traffic.

The following ports must use system VLANs:

- Control and packet VLANs in the uplinks that communicate with the VSM.
- The Management VLAN in the uplinks and port profiles (that is, the Ethernet and vEthernet ports) and VMware kernel NICs used for VMware vCenter Server connectivity, Secure Shell (SSH), or Telnet connections.
- The VLAN that is used for remote storage access (iSCSI or NFS).

**Caution**

You must use system VLANs sparingly and only as described in this section. Only 32 system port profiles are supported.

After a system port profile has been applied to one or more ports, you can add more system VLANs, but you can only delete a system VLAN after you remove the port profile from service. This action prevents you from accidentally deleting a critical VLAN, such as a host management VLAN or a VSM storage VLAN.

**Note**

One VLAN can be a system VLAN on one port and a regular VLAN on another port in the same ESX host.

To delete a system VLAN, see the *Cisco Nexus 1000V Port Profile Configuration Guide*.

Installation Overview

Installing the Cisco Nexus 1000V Manually


When you install the Cisco Nexus 1000V manually, you download and install all of the necessary software. This installation method gives you the option of deploying Layer 2 or Layer 3 connectivity between the VSM and VEMs. Layer 3 connectivity is the preferred method. For an example of the Layer 3 installation topology, see [Topology for Layer 3 Control Mode, on page 6](#). If you want to use Layer 2 connectivity, see [Topology for Layer 2 Control Mode, on page 7](#).

Recommended Topologies

Topology for Layer 3 Control Mode

Layer 3 control mode is the preferred method of communication between the VSM and VEMs. This figure shows an example of a Layer 3 control mode topology where redundant VSM VMs are installed. The software for the primary VSM is installed on ESXi 1, and the software for the secondary VSM is installed on ESXi 2.

Figure 3: Layer 3 Control Mode Topology Diagram



352064


Topology for Layer 2 Control Mode

**Note**

Layer 3 control mode is the preferred method for communications between the VSM and the VEMs. For a topology diagram, see [Topology for Layer 3 Control Mode](#), on page 6.

In Layer 2 control mode, the VSM and VEMs are in the same subnet. This figure shows an example of redundant VSM VMs, where the software for the primary VSM is installed on ESXi 1, and the software for the secondary VSM is installed on ESXi 2.

Figure 4: Layer 2 Control Mode Topology Diagram



Redundant Cisco Nexus 1000V VSMs
Primary and secondary VSMs form an HA Pair

- Management — VLAN 260, vmnic 0
- Control — VLAN 260, vmnic 0
- Packet — VLAN 260, vmnic 0
- Data — VLAN 20, vmnic 1

2099903

Control and Management VLAN Topology Options

You can deploy the control and management interfaces on separate VLANs or on the same VLAN.

