



Overview

This chapter contains the following sections:

- [Information about Layer 2 Switching, page 1](#)
- [Layer 2 Ethernet Switching, page 4](#)
- [MAC Address Tables, page 4](#)
- [VLANs, page 5](#)
- [Private VLANs, page 5](#)
- [IGMP Snooping, page 5](#)

Information about Layer 2 Switching

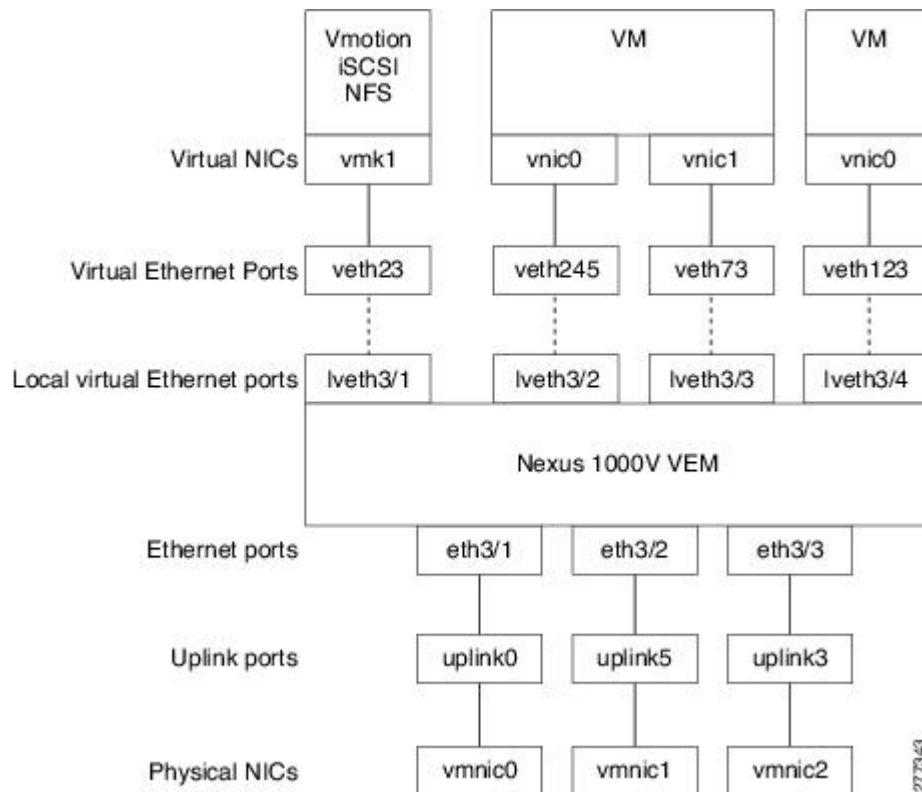
VEM Port Model

The Cisco Nexus 1000V differentiates the following Virtual Ethernet Module (VEM) ports:

- VEM Virtual Ports
- VEM Physical Ports

The following figure shows how VEM ports are bound to physical and virtual VMware ports.

Figure 1: VEM Port View



VEM Virtual Ports

The virtual side of the VEM maps together the following three layers of ports:

Virtual NICs

There are types of Virtual NICs in VMware. The virtual NIC (vnic) is part of the VM, and represents the physical port of the host which is plugged into the switch. The virtual kernel NIC (vmknic) is used by the hypervisor for management, iSCSI, NFS and other network access needed by the kernel. The vswif (not shown) appears only in COS-based systems, and is used as the VMware management port. Each of these types maps to a veth port within Cisco Nexus 1000V.

Virtual Ethernet Ports

A virtual Ethernet port (vEth) represents a port on the Cisco Nexus 1000V Distributed Virtual Switch. Cisco Nexus 1000V has a flat space of vEth ports, 0...n. These vEth ports are what the virtual “cable” plugs into, and are moved to the host that the VM is running on. Virtual Ethernet ports are assigned to port groups.

VEM Physical Ports

The physical side of the VEM includes the following from top to bottom:

Uplink Ports

Each uplink port on the host represents a physical interface.

Ethernet Ports

Each physical port added to Cisco Nexus 1000V appears as a physical Ethernet port, just as it would on a hardware-based switch.



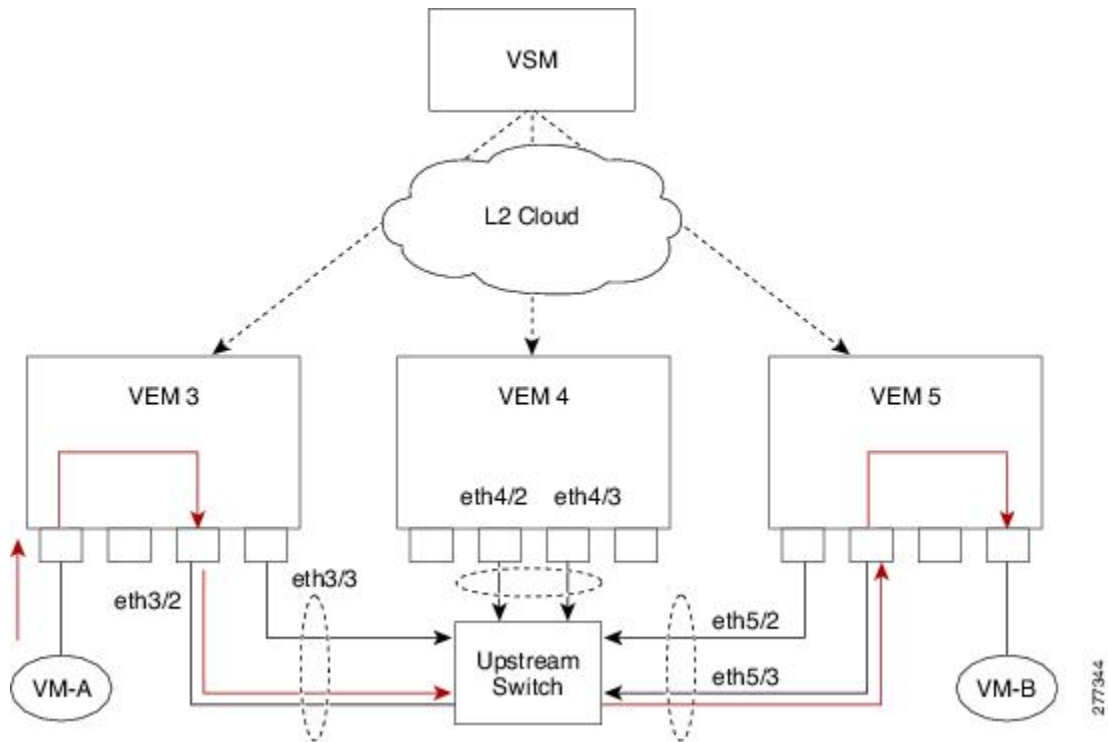
Note

There is no fixed relationship between the uplink number and number, and these can be different on different hosts, and can change throughout the life of the host.

VSM Port Model

The following figure shows the VSM view of the network.

Figure 2: VSM View



The Virtual Supervisor Module (VSM) has the following ports or interfaces:

Virtual Ethernet Interfaces

Virtual Ethernet interfaces (vEths) can be associated with any of the following:

- A virtual machine vNIC on the ESX host
- A virtual machine kernel NIC on the ESX host
- A virtual switch interface on an ESX COS host

Physical Ethernet Interfaces

Physical Ethernet interfaces (Eths) correspond to the physical NICs on the ESX host.

Port Channel Interfaces

The physical NICs of an ESX host can be bundled into a logical interface called a port channel interface.

Switching Traffic Between VEMs

Each VEM attached to the VSM forwards traffic to and from the server as an independent and intelligent line card. Each VLAN uses its forwarding table to learn and store MAC addresses for ports connected to the VEM.

The following figure shows the traffic flow between two VMs on different VEMs.

Layer 2 Ethernet Switching

The congestion related to high bandwidth and large numbers of users can be solved by assigning each device (for example, a server) to its own 10-, 100-, 1000-Mbps, or 10-Gigabit collision domain. Because each LAN port connects to a separate Ethernet collision domain, servers in a switched environment realize full bandwidth access.

Full duplex allows two stations to transmit and receive at the same time. This is unlike 10/100-Mbps Ethernet, which usually operates in half-duplex mode, so that stations can either receive or transmit but not both. When packets can flow in both directions simultaneously, the effective Ethernet bandwidth doubles. 1/10-Gigabit Ethernet operates in full-duplex only.

Each LAN port can connect to a single workstation or server or to another device through which workstations or servers connect to the network.

To reduce signal degradation, each LAN port is considered to be an individual segment. When stations connected to different LAN ports need to communicate, frames are forwarded from one LAN port to the other at wire speed to ensure full bandwidth for each session.

MAC Address Tables

To switch frames between LAN ports efficiently, a MAC address table is maintained. The MAC address of the sending network is associated with the LAN port on which it was received.

For more information about MAC address tables, see [Configuring the MAC Address Table](#).

VLANs

A VLAN is a switched network that is logically segmented by function, project team, or application, without regard to the physical locations of the users. VLANs have the same attributes of physical LANs, but you can group end stations even if they are not physically located on the same LAN segment.

Any switchport can belong to a VLAN, and unicast, broadcast, and multicast packets are forwarded and flooded only to end stations in that VLAN. Each VLAN is considered a logical network, and packets destined for stations that do not belong to the VLAN must be forwarded through a bridge or a router.

All ports, including the management port, are assigned to the default VLAN (VLAN1) when the device first comes up.

Up to 4094 VLANs are supported in accordance with the IEEE 802.1Q standard. These VLANs are organized into several ranges for different uses. Some of these VLANs are reserved for internal use by the device and are not available for configuration.

**Note**

Inter-Switch Link (ISL) trunking is not supported on the Cisco Nexus 1000V.

See [Configuring VLANs](#) for information about VLAN numbering and configuring VLANs.

Private VLANs

Private VLANs (PVLANS) are used to segregate Layer 2 ISP traffic and convey it to a single router interface. PVLANS achieve device isolation by applying Layer 2 forwarding constraints that allow end devices to share the same IP subnet while being Layer 2 isolated. In turn, the use of larger subnets reduces address management overhead.

See [Configuring Private VLANs](#) for more information.

IGMP Snooping

The Internet Group Management Protocol (IGMP) snooping software examines Layer 2 IP multicast traffic within a VLAN to discover the ports where interested receivers reside. Using the port information, IGMP snooping can reduce bandwidth consumption in a multi-access LAN environment to avoid flooding the entire VLAN. The IGMP snooping feature tracks which ports are attached to multicast-capable routers to help the routers forward IGMP membership reports. The IGMP snooping software responds to topology change notifications. By default, IGMP snooping is enabled on the device.

See [Configuring IGMP Snooping](#) for more information.

