Overview

This chapter contains the following sections:

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 the VEM view of the network.

VEM Virtual Ports

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

Virtual NICs

There are two types of virtual NICs (vNICs). One vNIC represents a network interface on a Virtual Machine (VM), which emulates a physical port for the virtual host. The other vNIC is an internal port used by the hypervisor for management, iSCSI, and other network access. Each of these vNICs maps to a Virtual Ethernet port within the Cisco Nexus 1000V.

Virtual Ethernet Ports

A virtual Ethernet port (vEth) represents a port on the Cisco Nexus 1000V Distributed Virtual Switch. The Cisco Nexus 1000V has a flat space of vEth ports, 1...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 profiles.

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

Note There is no initial relationship between a physical NIC assigned to the virtual distributed switch and the resulting ethernet port created on the VSM. However, after this is done, the PNIC and ethernet port mapping remains consistent as long at the port continues to be mapped to the virtual distributed switch, even when the physical host is rebooted.

VSM Port Model

The following figure shows the VSM view of the network.

Figure 1: VSM View

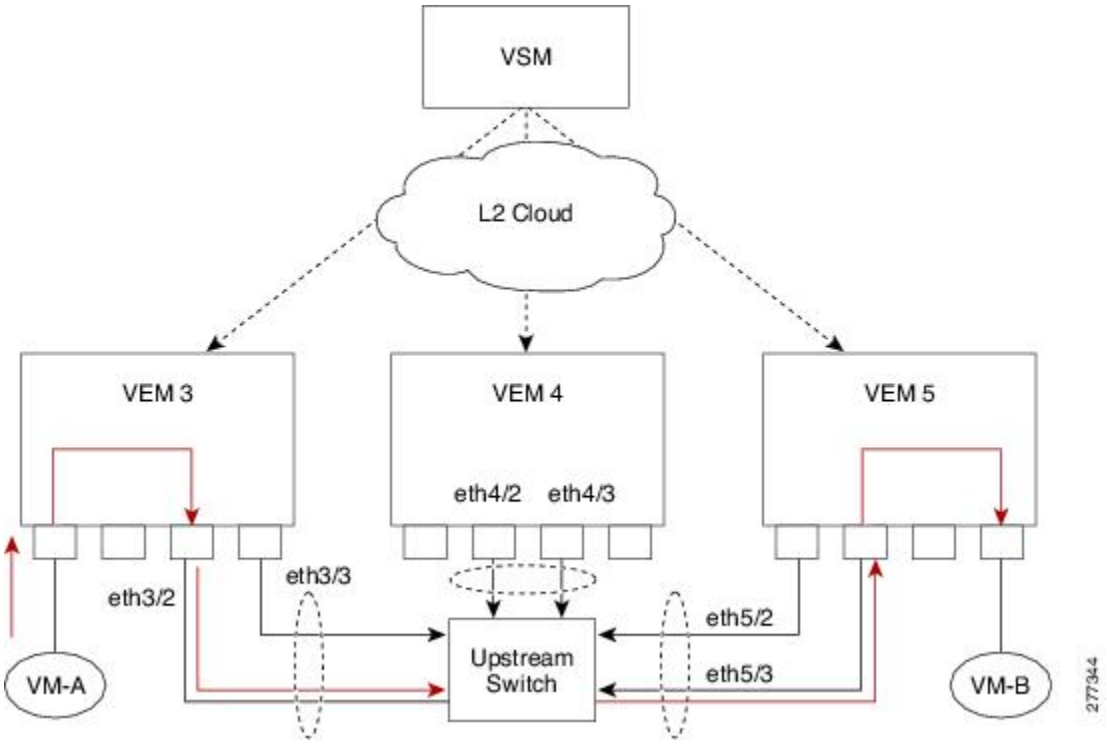

The Virtual Supervisor Module (VSM) has the following ports or interfaces:

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 that is 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 that are connected to the VEM.

See the following figure to see how traffic flows between 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. 10/100-Mbps Ethernet 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 mode 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.

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.

You can configure VLANs using the Cisco Nexus 1000V for KVM CLI, OpenStack Horizon Dashboard, or the OpenStack CLI. Although you can continue to configure VLANs using the Cisco Nexus 1000V for KVM CLI, Cisco recommends that you configure VLANs using OpenStack. For information about OpenStack, see the *Cisco Nexus 1000V for KVM Virtual Network Configuration Guide*.

Note You configure VLANs as VM subnets through OpenStack.

You must consistently use OpenStack for all VM network and subnet configuration. If you use *both* OpenStack and the VSM to configure VM networks and subnets, the OpenStack and the VSM configurations can become out-of-sync and result in faulty or inoperable network deployments.

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

Control VLANs

A control VLAN is used for communication between the VSM and the VEMs within a switch domain. The control interface is the first interface on the VSM.

A control VLAN is used for the following:

- VSM configuration commands to each VEM and their responses.
- VEM notifications to the VSM. For example, a VEM notifies the VSM of the attachment or detachment of ports to the Distributed Virtual Switch (DVS).
- VEM NetFlow exports that are sent to the VSM, where they are forwarded to a NetFlow Collector.
- VSM active to standby synchronization for high availability.

Management VLANs

A management VLAN, which is used for system login and configuration, corresponds to the mgmt0 interface. The mgmt0 interface appears as the mgmt0 port on a Cisco switch and is assigned an IP address (IPv4 or IPv6). Although the management interface is not used to exchange data between the VSM and VEM, it is used to establish and maintain the connection between the VSM and VMware vCenter Server in Layer 2 mode. In Layer 3 mode (default), when the mgmt0 interface (default) is used for Layer 3 connectivity on the VSM, the management interface communicates with the VEMs and the VMware vCenter Server.

The management interface is the second interface on the VSM.

Packet VLANs

Note A packet VLAN is not a component of the Layer 3 control mode. If you are using Layer 3 control mode, you do not need a packet VLAN.

Similar to the control VLAN, a packet VLAN is used for communication between the VSM and the VEMs within a switch domain.

A packet VLAN is used to tunnel network protocol packets between the VSM and the VEMs such as the Cisco Discovery Protocol (CDP), Link Aggregation Control Protocol (LACP), and Internet Group Management Protocol (IGMP).

The packet interface is the third interface on the VSM.

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. The use of larger subnets reduces address management overhead.

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.