

Cisco Nexus 1000V Switch for Microsoft Hyper-V

Version 1

Deployment Guide

July 2013

<u>What You Will Learn</u>	4
<u>Audience</u>	4
<u>Overview</u>	4
<u>Cisco Nexus 1000V Switch Components</u>	4
<u>Template-Based Network Policy</u>	5
<u>Microsoft Hyper-V Networking</u>	6
<u>Microsoft SCVMM Networking</u>	7
<u>Switch Extension Managers</u>	7
<u>Logical Switch</u>	8
<u>Port Classifications and Virtual Machine Networks</u>	10
<u>Port Classifications</u>	10
<u>Logical Network and Network Sites</u>	12
<u>Virtual Machine Networks</u>	13
<u>IP Address Pools</u>	16
<u>Cisco Nexus 1000V Switch Concepts and Components</u>	16
<u>Cisco NX-OS Software</u>	16
<u>Virtual Supervisor Module</u>	17
<u>VSM Interfaces</u>	17
<u>Control Interface</u>	18
<u>Management Interface</u>	18
<u>Packet Interface</u>	18
<u>Virtual Ethernet Module</u>	18
<u>VSM-to-VEM Communication</u>	19
<u>Cisco Nexus 1000V Switch Installation</u>	24
<u>Cisco Nexus 1000V Switch Features</u>	25
<u>Switch Forwarding</u>	25
<u>MAC Address Learning</u>	25
<u>Loop Prevention</u>	26
<u>Switch Port Interfaces</u>	27
<u>Opaque Data</u>	27
<u>Port Profiles</u>	28
<u>Virtual Ethernet Port Profiles</u>	28
<u>Ethernet Port Profiles</u>	29
<u>Network Segments</u>	30
<u>Dynamic Port Profiles</u>	31
<u>System Network Segments</u>	31
<u>Policy Mobility</u>	32
<u>Licensing</u>	32
<u>Cisco Nexus 1000V Switch Network Design</u>	33
<u>Deployment Topology 1: Recommended</u>	33
<u>Management Microsoft Hyper-V Clusters and Hosts</u>	33
<u>VSM High Availability</u>	34
<u>Layer 3 Mode</u>	34
<u>Data Virtual Machine Cluster</u>	34
<u>Management, Live Migration, and Cluster Traffic</u>	34
<u>Deployment Topology 2: For Customers with Limited pNICs per Microsoft Hyper-V Host</u>	35

<u>Deploying Cisco Nexus 1000V on Cisco Unified Computing System</u>	35
<u>Cisco Virtual Interface Card</u>	36
<u>Service Profile Design</u>	36
<u>Quality of Service</u>	38
<u>Upstream Switch Connectivity</u>	39
<u>Standard PortChannel</u>	40
<u>LACP Offload</u>	40
<u>Special PortChannel</u>	41
<u>vPC-HM MAC Address Pinning</u>	41
<u>vPC-HM Subgroups</u>	43
<u>Load Balancing</u>	43
<u>Source-Based Hashing</u>	44
<u>Flow-Based Hashing</u>	44
<u>Network-State Tracking</u>	44
<u>Cisco Nexus 1000V Switch for Microsoft Hyper-V Sample Configuration</u>	45
<u>Conclusion</u>	47
<u>For More Information</u>	47

What You Will Learn

This document provides design and configuration guidance for deployment of the Cisco Nexus® 1000V Switch for Microsoft Hyper-V. For detailed installation and configuration documentation, refer to the respective Cisco® documentation at <http://www.cisco.com/en/US/products/ps13056/index.html>.

Audience

This document is intended for network architects, network engineers, virtualization administrators, and server administrators interested in understanding the deployment of the Cisco Nexus 1000V Switch for Microsoft Hyper-V.

Overview

When server virtualization is implemented, the edge of the network is pushed from the traditional location in the network access layer, implemented in physical switches, to the virtual network access layer that is implemented in software in the server hypervisor. The Cisco Nexus 1000V Switch is an intelligent virtual network access layer switch that runs Cisco NX-OS Software, Cisco's data center operating system that runs on all Cisco data center products. Operating inside the Microsoft Hyper-V hypervisor, the Cisco Nexus 1000V supports Cisco Virtual Network Link (VN-Link) server virtualization technology to provide:

- Policy-based virtual machine connectivity
- Mobile virtual machine security and network policy
- Nondisruptive operating model for your server virtualization and networking teams

When server virtualization is implemented in the data center, servers and virtual machines are not managed the same way as physical servers. Server virtualization is treated as a special deployment, leading to longer deployment time, with more coordination needed among server, network, storage, and security administrators. With the Cisco Nexus 1000V, you have a consistent networking feature set and configuration and provisioning model for both the physical and the virtual networks. Virtual machine networks can use the same network configuration, security policy, diagnostic tools, and operating models as physical server deployments that are connected to physical switches. This unified approach provides faster deployment and troubleshooting and makes the administration of virtualization environments essentially the same as for nonvirtualized deployments.

Developed in close collaboration with Microsoft, the Cisco Nexus 1000V Switch is certified by Microsoft and integrates with Microsoft Windows Server and Microsoft System Center Virtual Machine Manager (SCVMM). You can use the Cisco Nexus 1000V to manage your virtual machine connectivity with confidence in the integrity of the server virtualization infrastructure.

Cisco Nexus 1000V Switch Components

Cisco Nexus 1000V Switches have two main components:

- Virtual supervisor module (VSM)
- Virtual Ethernet module (VEM)

The VSM provides the switch control and management plane, and the VEM provides the data plane for the switch (Figure 1). The VSM can run as a virtual machine on any Microsoft Hyper-V host or as a virtual service node on the Cisco Nexus 1010 and 1110. The VEM runs as a plug-in (extension) to the Microsoft Hyper-V switch in the hypervisor kernel, providing switching between virtual machines.

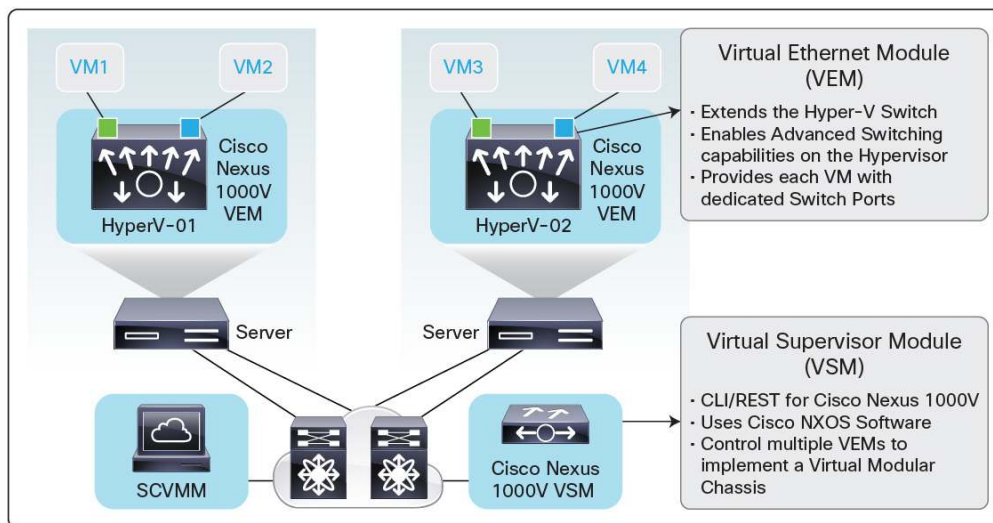
Cisco Nexus 1000V sees the VSMs and VEMs as modules. In the current release, a single VSM can manage up to 64 VEMs. The VSMs are always associated with slot numbers 1 and 2 in the virtual chassis. The VEMs are sequentially assigned to slots 3 through 66 based on the order in which their respective hosts were added to the Cisco Nexus 1000V Switch.

The Cisco Nexus 1000V Switch provides Layer 2 switching, advanced networking functions (security and monitoring), and a common network management model for Microsoft Hyper-V servers by implementing a **forwarding extension** for the Microsoft Hyper-V Extensible Switch. For more information about Microsoft's Extensible Switch architecture, refer to the Microsoft documentation at <http://technet.microsoft.com/en-us/library/hh831452.aspx>.

Microsoft strongly recommends the use of Microsoft SCVMM as a centralized platform for managing Compute, Network and Storage resources in the data center. Cisco Nexus 1000V integrates with Microsoft SCVMM 2012 SP1 to allow the virtualization administrator to deploy virtual machines to the virtual access layer. The server and virtualization teams must deploy Microsoft SCVMM to use Cisco Nexus 1000V to manage virtual networking.

Note: Cisco Nexus 1000V integrates with Microsoft SCVMM SP1 UR2 Version 3.1.6020.0 and later. Verify that you are running the minimum required Microsoft SCVMM version before installing Cisco Nexus 1000V.

Figure 1. Cisco Nexus 1000V Switch Components



Template-Based Network Policy

A unique aspect of the Cisco Nexus 1000V is the way that network policy is defined and deployed. Today, when an application is deployed on a physical server, the upstream network switches are configured. VLANs and network policies are applied to the physical switches to help ensure that the application and host run securely and are highly available. For Cisco switches, this management model requires administrators to enter configuration mode and apply a series of switch commands that define the interface configuration.

Often, identical configurations are needed on multiple interfaces of the same switch or on different switches. This management model requires server administrators to depend on network administrators to reconfigure the network each time a server is brought online. This process can create unnecessary delays in deployment of new

servers. When this model is extended to virtual machines, which are deployed on a larger scale than physical hosts, it leads to longer deployment times.

The Cisco Nexus 1000V provides an excellent and improved management and configuration model to manage virtual switches and virtual network policy. In this model, network administrators define a network policy template that virtualization or server administrators can apply to all virtual machines that require the same network policy. These policy templates are referred to as port profiles.

Port profiles create a unique collaborative model, giving server administrators the autonomy to provision new virtual machines without waiting for network reconfigurations to be implemented in the physical network infrastructure. For network administrators, the combination of the Cisco Nexus 1000V feature set and the capability to define a port profile using the same syntax as for existing physical Cisco switches helps ensure that consistent policy is enforced without the burden of having to manage individual virtual switch ports.

In addition to reducing deployment time, port profiles help maintain separation of duties for the server and the network administrators. Although virtual switches are created on Microsoft Hyper-V servers, port profiles defined on the VSM enable the Cisco Nexus 1000V administrator to maintain and enforce consistent network policy on the virtual access layer. In this model, the server administrator is not burdened with implementing virtual network policy. Instead, the administrator simply applies the correct port classification to each virtual machine deployed on the Cisco Nexus 1000V logical switch.

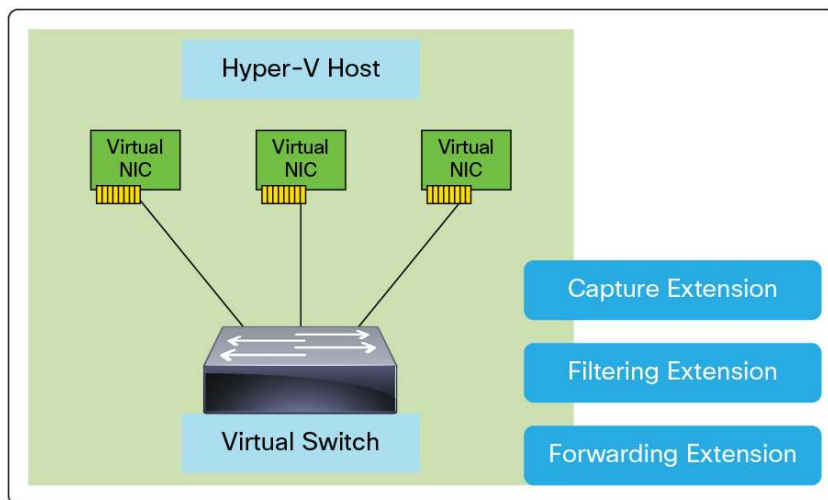
Microsoft Hyper-V Networking

To understand the Cisco Nexus 1000V Switch for Microsoft Hyper-V, you must first understand the basics of Microsoft Hyper-V 2012 networking. In Microsoft Windows Server 2012, Microsoft redesigned the native virtual switch that is shipped with Microsoft Hyper-V to implement an extensible switch. The new Microsoft Hyper-V extensible switch architecture allows third-party vendors to provide extensions that enhance the capabilities of the native switch supported on Microsoft Hyper-V (Figure 2). The extensions supported are:

- Filter: Provides the capability to filter certain flows
- Capture: Provides the capability to capture and redirect certain flows
- Forward: Provides the capability to filter and capture flows and perform Layer 2 forwarding

The Cisco Nexus 1000V VSM implements a forwarding extension to enhance the capabilities of the Microsoft Hyper-V extensible switch. For more information about the Microsoft Hyper-V extensible switch, refer to the Microsoft documentation at [http://msdn.microsoft.com/en-us/library/windows/hardware/hh582268\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/hardware/hh582268(v=vs.85).aspx).

Figure 2. Microsoft Hyper-V Extensible Switch



Microsoft SCVMM Networking

This section introduces some Microsoft SCVMM concepts that are relevant to Cisco Nexus 1000V and its integration with Microsoft SCVMM. For detailed information about Microsoft SCVMM networking, please refer to the Microsoft documentation.

Switch Extension Managers

Microsoft SCVMM views third-party applications such as the Cisco Nexus 1000V VSM as configuration providers. Before the Cisco Nexus 1000V VSM is used as a forwarding extension in a Microsoft Hyper-V switch, the VSM must be registered with Microsoft SCVMM as a virtual switch extension manager (VSEM). After the VSM is registered as a VSEM, Microsoft SCVMM imports all configurations defined by the VSM administrator on the Cisco Nexus 1000V.

A VSEM is created by connecting to the VSM management IP address and the switch administrator credentials. In Figure 3, a Cisco Nexus 1000V VSM is being added as a VSEM by connecting to the switch management IP address of 10.10.1.10 using HTTP. A RunAs account called VSM Admin has been created using the switch administrator credentials.

Figure 3. Adding the Cisco Nexus 1000V as Virtual Switch Extension Manager

The screenshot shows a Windows-style wizard window titled "Add Virtual Switch Extension Manager Wizard". The "General" tab is selected in the left-hand pane, which also lists "Host Groups" and "Summary". The main area contains the following fields and instructions:

- Enter connection settings for the extension manager to add**
- Select a manufacturer, model, and configuration provider for the extension manager. Enter the connection string and credentials to be used.**
- Manufacturer:** Cisco Systems (dropdown menu)
- Model:** Nexus 1000V (dropdown menu)
- Provider:** Cisco Systems Nexus 1000V (dropdown menu)
- Connection string:** http://10.10.1.10 (text box)
- RunAs account:** VSM Admin (text box) with a "Browse..." button next to it.

At the bottom right, there are three buttons: "Previous", "Next" (highlighted in blue), and "Cancel".

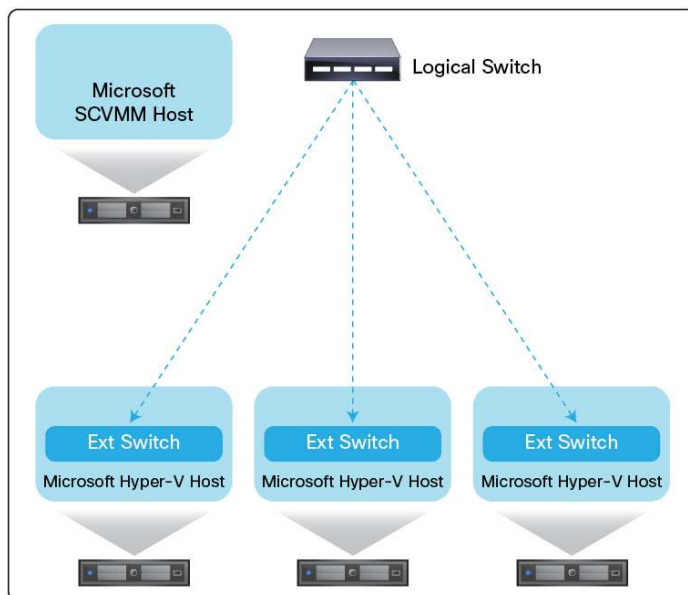
Note: Before the Cisco Nexus 1000V VSM is registered as a VSEM, the Cisco provider Microsoft installation (MSI) file, which is available as part of the download bundle, must be installed on the Microsoft SCVMM server. The provider MSI file enhances Microsoft SCVMM so that it communicates with the Cisco Nexus 1000V VSM. This file is automatically installed on the Microsoft SCVMM server host when the installer application is used.

Logical Switch

Microsoft introduced the concept of a Logical Switch in Microsoft SCVMM 2012 SP1 to help ensure consistent virtual network policy across Microsoft Hyper-V servers. A Logical Switch is a switch template created on Microsoft SCVMM to help ensure that consistent virtual network policy is applied to multiple Microsoft Hyper-V hosts. Configuration changes made to the template immediately are reflected on all the hosts on which the Logical Switch instance has been created.

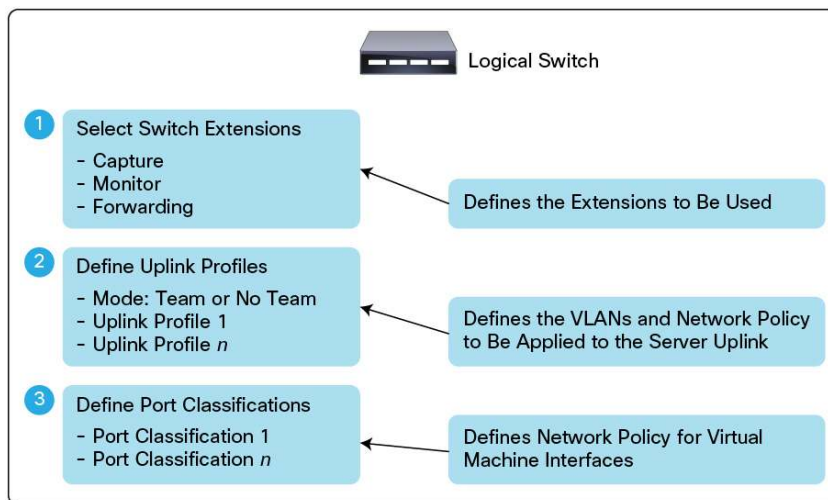
When Cisco Nexus 1000V is used with Microsoft SCVMM, a Logical Switch that uses the Cisco Nexus 1000V as a forwarding extension is created on Microsoft SCVMM. This Logical Switch is then instantiated on all Microsoft Hyper-V hosts on which virtual networking needs to be managed with Cisco Nexus 1000V (Figure 4). As a result of this process, a Microsoft Hyper-V extensible switch with the Cisco Nexus 1000V as a forwarding extension is created on Microsoft Hyper-V hosts. In addition to providing consistent network policy across multiple Microsoft Hyper-V hosts, the Cisco Nexus 1000V implementation of a distributed virtual switch gives the administrator a single view of the virtual network from the VSM console.

Figure 4. Logical Switch Created on Microsoft SCVMM and Instantiated on Microsoft Hyper-V Hosts



A Logical Switch template definition contains the information shown in Figure 5.

Figure 5. Logical Switch Defines a Set of Uplink Profiles and Port Classifications



Uplink profiles and port classifications are explained in the next sections of this document.

Note: When a Cisco Nexus 1000V Logical Switch is created on Microsoft SCVMM, only one extension is used. The Cisco Nexus 1000V is used as a forwarding extension.

Note: In this release, only one Cisco Nexus 1000V Logical Switch can be created per Microsoft Hyper-V host: that is, creation of more than one VEM per Microsoft Hyper-V host is not supported.

Port Classifications and Virtual Machine Networks

A virtual machine is deployed on a Logical Switch by specifying a specific port classification and a virtual machine network configured on the Logical Switch. The port classification identifies all the network policy that must be applied to a virtual machine's virtual Ethernet (vEth) interface, and the virtual machine network identifies the network on which the virtual machine is deployed.

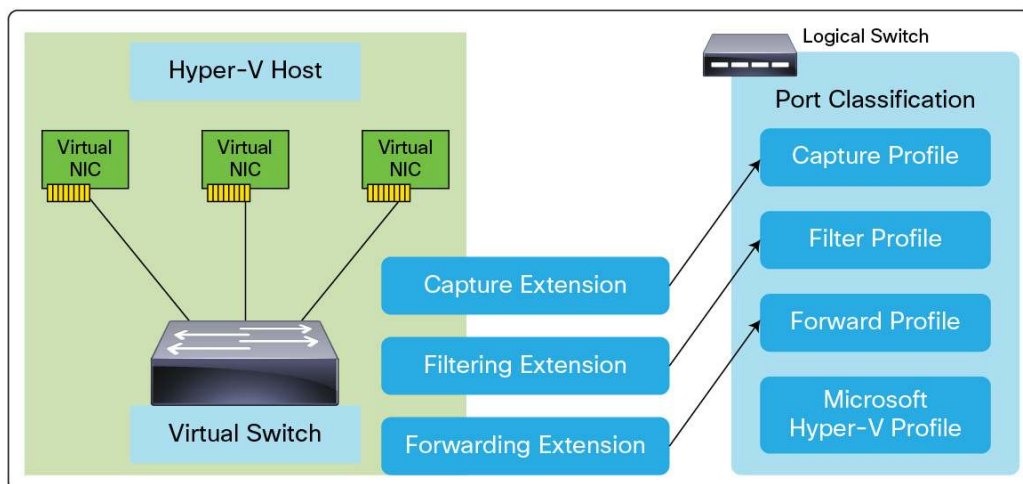
When the Cisco Nexus 1000V is used to manage the virtual access layer on Microsoft Hyper-V servers, the VSM administrator creates port profiles and network segments. The Microsoft SCVMM administrator uses the port profile created on the Cisco Nexus 1000V Switch for Microsoft Hyper-V to create a port classification. Similarly, the administrator uses a network segment created on the VSM to create a virtual machine network on Microsoft SCVMM.

Port Classifications

A port classification is a container or template of virtual network policy that is applied to a virtual machine's Ethernet interface. Microsoft SCVMM uses port classifications to define virtual network policy once and then apply the same policy to all virtual machines interfaces requiring identical configuration. Port classifications also help ensure that consistent network policy is applied to all virtual machine Ethernet interfaces. A change made to the port classification is immediately applied to all interfaces on which the classification is applied.

Microsoft Hyper-V extensible switch architecture allows multiple vendors to provide extensions that enhance the native switch's behavior. Each of these extensions can provide network policy to Microsoft SCVMM in the form of a port profile. A port classification combines network policy from the capture, filter, and forwarding extensions and the native switch to define all the policy that is to be applied to a virtual machine's Ethernet interface (Figure 6).

Figure 6. Port Classification: Collection of Port Profiles



Because the Cisco Nexus 1000V implements a forwarding extension and can perform capture and filter functions, in practice a port classification on a Cisco Nexus 1000V Logical Switch contains only a forwarding port profile.

In Figure 7, the Cisco Nexus 1000V administrator has defined a simple port profile called RestrictedProfile that applies an access control list (ACL) network policy.

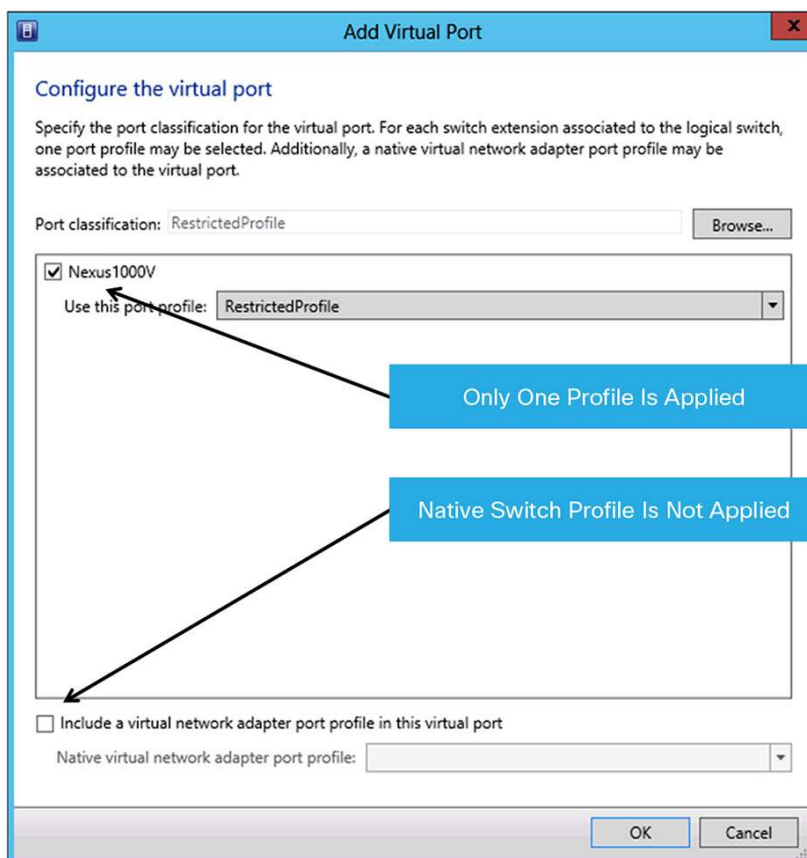
Figure 7. Simple Port Profile Defined on the Cisco Nexus 1000 VSM

```
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# ip access-list Restrict-ContractorVM
Nexus1000V(config-acl)# 10 deny ip 20.20.20.22/32 20.20.20.20/32
Nexus1000V(config-acl)# 20 permit ip any any
Nexus1000V(config-acl)# exit
Nexus1000V(config)#

Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# port-profile type vethernet RestrictedProfile
Nexus1000V(config-port-prof)# ip port access-group Restrict-ContractorVM in
Nexus1000V(config-port-prof)# no shutdown
Nexus1000V(config-port-prof)# state enabled
Nexus1000V(config-port-prof)# publish port-profile
Nexus1000V(config-port-prof)# exit
Nexus1000V(config)#
Nexus1000V(config)#
```

The Microsoft SCVMM administrator uses RestrictedProfile when he creates a port classification. In Figure 8, the administrator is creating a port classification, also called RestrictedProfile, with only one port profile: the RestrictedProfile port profile defined on the VSM.

Figure 8. Creating a Port Classification Using the Cisco Nexus 1000V Port Profile



Note: Port classifications are similar to port groups defined in VMware vCenter for VMware ESX environments. However, in VMware vCenter, creation of a port profile on the Cisco Nexus 1000V results in the automatic creation of a port group, whereas in Microsoft SCVMM, the user has to manually create a port classification. The extra step is needed because a port classification can represent network policies from more than one provider.

Logical Network and Network Sites

The Microsoft SCVMM networking hierarchy or object model consists of the following new network constructs: Logical Networks, Network Sites, and Virtual Machine Networks (VM Networks). Microsoft SCVMM introduced this hierarchy to enable the consumer of network resources (server and tenant administrators) to automate the provisioning of virtual machine networks. Also, this new networking hierarchy enables the server administrator to request and consume network resources without needing to understand the underlying network implementation details.

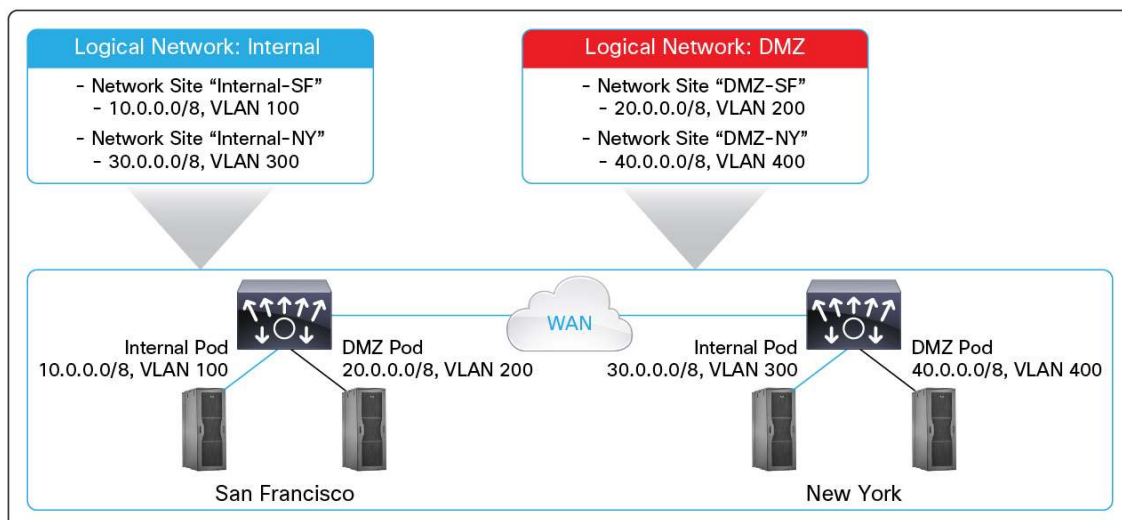
A Logical Network models distinct networks (or types of networks) managed by an enterprise. The Logical Network abstraction hides the VLANs and IP subnets that provide isolation from all the consumer and users (virtual machine network administrators, tenant administrators, and server administrators) other than the fabric administrator managing the physical fabric. In other words, a Logical Network is composed of one or more network sites, and each network site is a group of VLANs, IP subnets, or VLAN and IP subnet pairs.

Consider a hypothetical enterprise, Avatar.local, which operates two data centers: one in San Francisco and one in New York. The enterprise has two subnets at each site:

- The San Francisco site uses subnet 10.0.0.0/8 to provide an Internal network. The subnet 20.0.0.0/8 is used to create a DMZ network.
- The New York site uses subnet 30.0.0.0/8 to provide an Internal network. The subnet 40.0.0.0/8 is used to create a DMZ network.

To model the network fabric of Avatar.local, the Microsoft SCVMM fabric administrator creates two Logical Networks: Internal and DMZ. The internal Logical Network has two Network Sites: 10.0.0.0/8 and 30.0.0.0/8. The DMZ Logical Network has two Network Sites: 20.0.0.0/8 and 40.0.0.0/8 (Figure 9).

Figure 9. Logical Network Modeling the Physical Network Fabric



When the Cisco Nexus 1000V is used to manage the virtual access layer on Microsoft Hyper-V, Logical Networks and Network Sites are created from the VSM. Network sites are referred to as network segment pools on the VSM because they are a collection of VLAN and IP subnets: that is, network segments. Figure 10 shows an example of how a Logical Network and network segment pool (Network Site) are created on Microsoft SCVMM.

Figure 10. Creating a Logical Network and Network Segment Pool on Microsoft SCVMM

```
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# nsm logical network DMZ
Nexus1000V(config-logical-net)# description Logical Network for DMZ VMs
Nexus1000V(config-logical-net)# exit
Nexus1000V(config)#
Nexus1000V(config)#

Nexus1000V(config)#
Nexus1000V(config)# nsm network segment pool DMZ-SFO
Nexus1000V(config-net-seg-pool)# member-of logical network DMZ
Nexus1000V(config-net-seg-pool)# exit
Nexus1000V(config)# nsm network segment pool DMZ-NY
Nexus1000V(config-net-seg-pool)# member-of logical network DMZ
Nexus1000V(config-net-seg-pool)# exit
Nexus1000V(config)#
```

Virtual Machine Networks

A Virtual Machine Network or VM Network provides isolation to virtual machines deployed on the Microsoft Hyper-V virtual access layer.

A Virtual Machine Network of a specific type is created by the server administrator by requesting a network resource from the Logical Network, which is a pool of Network Sites or network segments. In the example in Figure 10, the fabric administrator has created two types of Logical Networks: DMZ and Internal. The server administrator can now create virtual machine networks of either the DMZ or the internal type.

A virtual machine network defines a routing domain similar to Cisco virtual routing and forwarding (VRF) instances and can contain one or more virtual subnets. Each virtual subnet defines a Layer 2 broadcast domain.

Microsoft SCVMM 2012 SP1 supports four types of Virtual Machine Network:

- VLAN
- Network virtualization using generic route encapsulation (NVGRE)
- External (VLAN)
- Not Isolation

In the first release of Cisco Nexus 1000V, only external VLAN-based isolation is supported. Other segmentation technologies such as NVGRE and VXLAN are not supported in the first release. VLAN-backed network segments defined on the Cisco Nexus 1000V can be used by the Microsoft SCVMM administrator to create external Virtual Machine Networks.

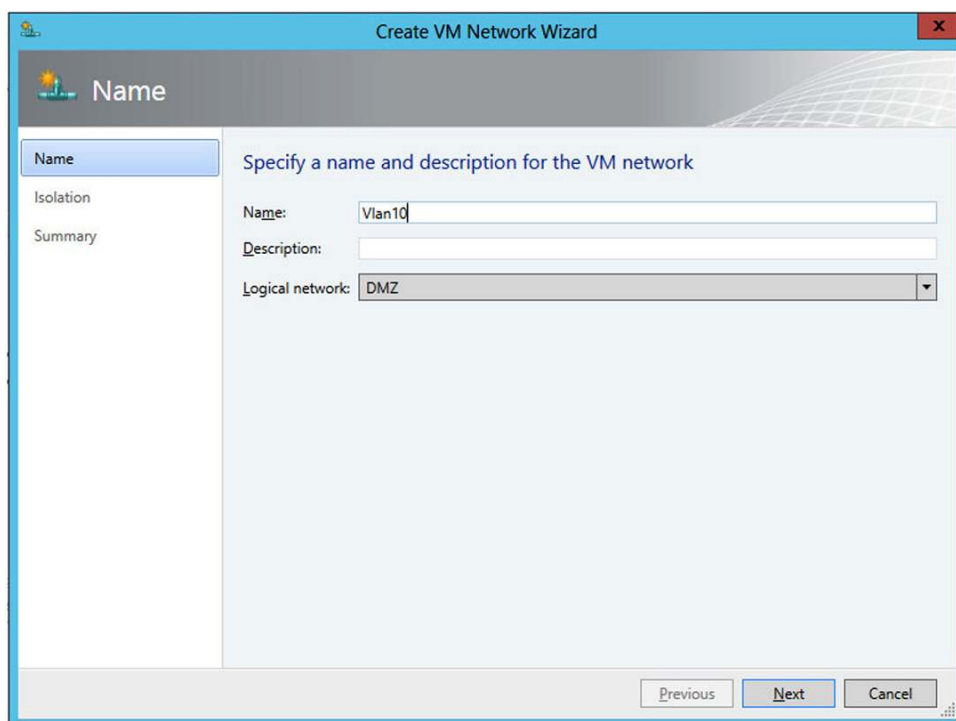
In Figure 11, the VSM administrator created a VLAN-backed network segment. The administrator is also applying the IP pool Vlan10IPPool to the Vlan10 network segment.

Figure 11. Creating a VLAN-Backed Network Segment

```
Nexus1000V(config)#
Nexus1000V(config)# nsm network segment Vlan10
Nexus1000V(config-net-seg)# member-of network segment pool DMZ-SFO
Nexus1000V(config-net-seg)# switchport access vlan 10
Nexus1000V(config-net-seg)# switchport mode access
Nexus1000V(config-net-seg)# ip pool import template Vlan10IPPool
Nexus1000V(config-net-seg)# publish network segment
Nexus1000V(config-net-seg)# exit
Nexus1000V(config)#
```

When the Microsoft SCVMM administrator creates a virtual machine network, the administrator can pick one of the available externally supplied virtual machine networks for isolation. In Figure 12, the Microsoft SCVMM administrator is creating a virtual machine network of type DMZ. The administrator is choosing the Vlan10 network segment defined by the VSM administrator.

Figure 12. Creating a Virtual Machine Network That Uses an External Virtual Machine Network for Isolation



The screenshot shows the 'Create VM Network Wizard' window. The 'Name' tab is active, displaying the following fields:

- Name:** Vlan10
- Description:** (empty field)
- Logical network:** DMZ (selected from a dropdown menu)

On the left side, there are three tabs: 'Name', 'Isolation', and 'Summary'. At the bottom right, there are three buttons: 'Previous', 'Next', and 'Cancel'.

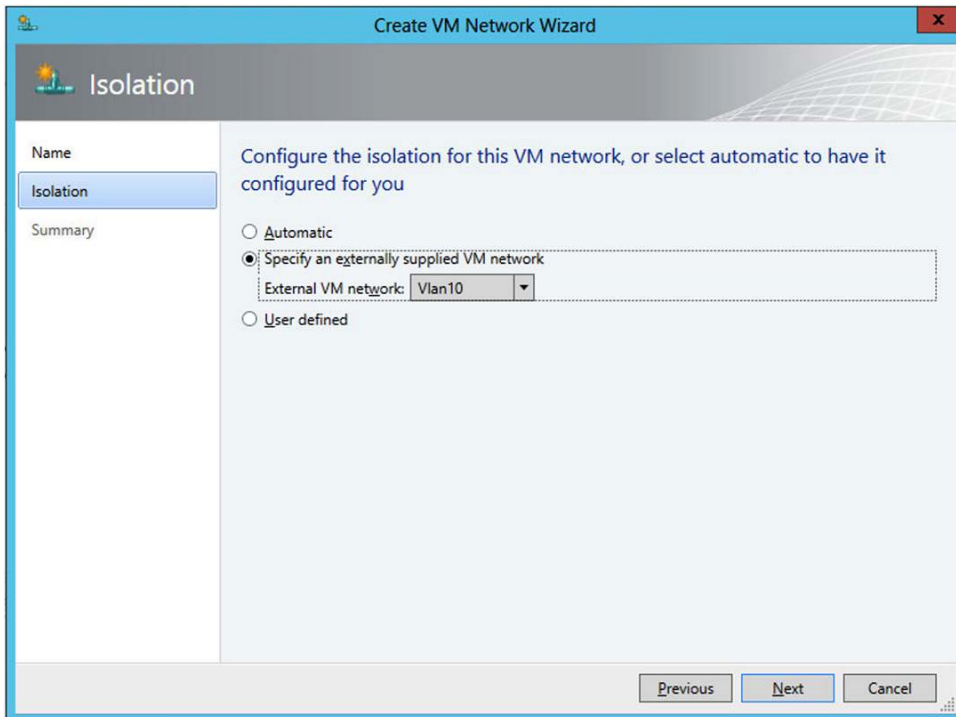
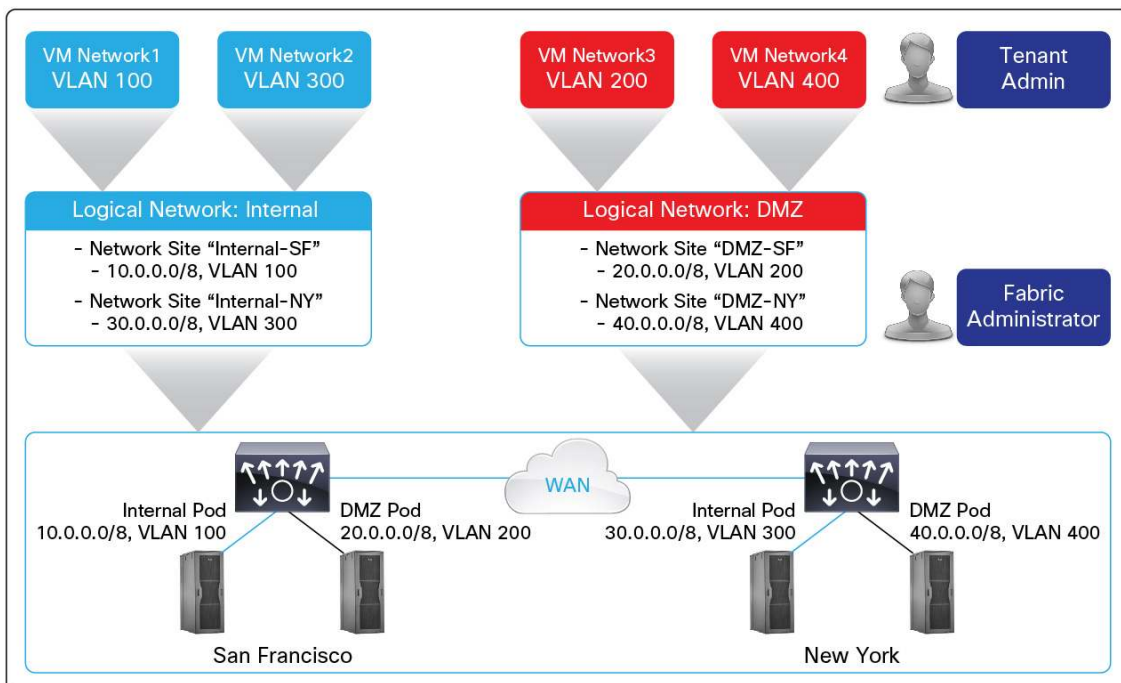


Figure 13 shows the Microsoft SCVMM network model of abstracting physical networks in terms of Logical Networks and then allowing the tenant administrator to create virtual machine networks using the resources available in Microsoft SCVMM Logical Networks.

Figure 13. Tenant Administrator Creating Virtual Machine Networks Using Resources Configured in Logical Networks



IP Address Pools

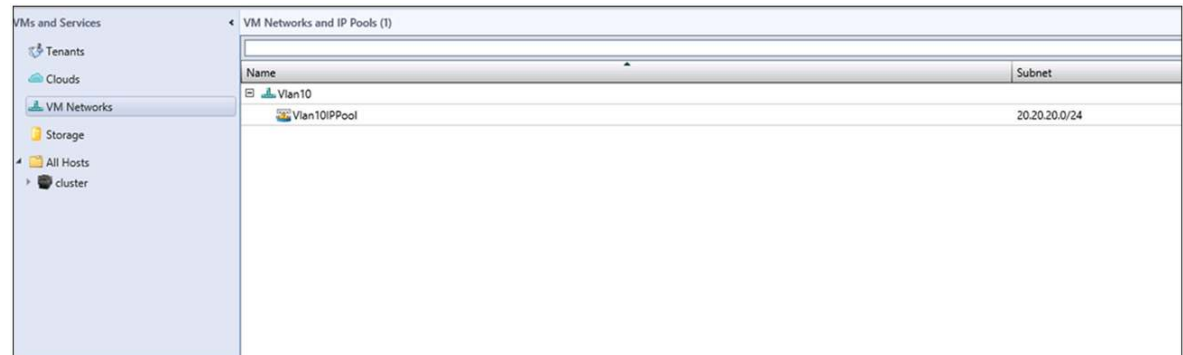
Microsoft SCVMM allows administrators to create IP address pools on virtual machine networks. When a virtual machine is deployed on a virtual machine network, an IP address from the pool is used as the static IP address on the virtual machine. When the Cisco Nexus 1000V is used to configure the Microsoft Hyper-V virtual network, the VSM administrator must define IP pools for a network segment. This network segment is then used by the Microsoft SCVMM administrator to create a Virtual Machine Network.

Figure 14 shows an example of IP pool creation on the VSM.

Figure 14. Virtual Machine Network with IP Pool in Microsoft SCVMM

```
Nexus1000V(config)#
Nexus1000V(config)# nsm ip pool template Vlan10IPPool
Nexus1000V(config-ip-pool-template)# ip address 20.20.20.50 20.20.20.70
Nexus1000V(config-ip-pool-template)# network 20.20.20.1 255.255.255.0
Nexus1000V(config-ip-pool-template)# default-router 20.20.20.1
Nexus1000V(config-ip-pool-template)# exit
Nexus1000V(config)#

Nexus1000V(config)#
Nexus1000V(config)# nsm network segment Vlan10
Nexus1000V(config-net-seg)# member-of network segment pool DMZ-SFO
Nexus1000V(config-net-seg)# switchport access vlan 10
Nexus1000V(config-net-seg)# switchport mode access
Nexus1000V(config-net-seg)# ip pool import template Vlan10IPPool
Nexus1000V(config-net-seg)# publish network segment
Nexus1000V(config-net-seg)# exit
Nexus1000V(config)#
```



Name	Subnet
Vlan10	
Vlan10IPPool	20.20.20.0/24

For more detailed information about creating virtual machine networks, please refer to the Cisco Nexus 1000V configuration guide.

Cisco Nexus 1000V Switch Concepts and Components

This section describes the major concepts and components of the Cisco Nexus 1000V Switches and the interaction of the components.

Cisco NX-OS Software

Cisco NX-OS is a data center-class operating system built with modularity, resiliency, and serviceability at its foundation. Running on all Cisco data center products, Cisco NX-OS helps ensure continuous availability and sets the standard for mission-critical data center environments. The self-healing and highly modular design of Cisco NX-OS makes zero-impact operations a reality and enables exceptional operation flexibility. Focused on the

requirements of the data center, Cisco NX-OS provides a robust and comprehensive feature set that can meet the Ethernet and storage networking requirements of present and future data centers.

With a command-line interface (CLI) like that of Cisco IOS® Software, Cisco NX-OS provides state-of-the-art implementations of relevant networking standards as well as a variety of true data center-class Cisco innovations.

Virtual Supervisor Module

The VSM provides the management and control plane functions for the Cisco Nexus 1000V Switches. Much like a supervisor module in a Cisco Nexus 7000 Series Switch, the VSM provides the switch control and management plane to the network administrator, coordinating configuration and functions across VEMs.

Unlike a traditional Cisco switch, in which the management plane is integrated into the hardware, on the Cisco Nexus 1000V the VSM is deployed either as a virtual machine on a Microsoft Hyper-V server or as a virtual service blade (VSB) on the Cisco Nexus 1010 or 1110 appliance (Figure 15).

Figure 15. VSM Representation



A high-availability deployment model of Cisco Nexus 1000V is usually recommended. Two VSMs are deployed in an active-standby configuration, with the first VSM functioning in the primary role and the other VSM functioning in a secondary role. If the primary VSM fails, the secondary VSM takes over.

Note that unlike in cross-bar-based modular switching platforms, the VSM is not in the data path. General data packets are not forwarded to the VSM to be processed, but are switched by the VEM directly. Hence, any disruption in VSM-to-VEM communication does not affect virtual machine traffic.

When the Cisco Nexus 1000V VSM is deployed on a Microsoft Hyper-V server, two installation options are available: automated installation using the installer application, and manual installation using the VSM virtual machine template on Microsoft SCVMM. When the VSM is deployed on the Cisco Nexus 1010 or 1110 device, the administrator must manually install the VSM VSB. For additional information about installation, please refer to the appropriate installation guides.

A VSM virtual machine must meet the following hardware requirements for the VSM to function correctly:

- 4 GB of virtual RAM (VRAM)
- 4 GB of hard disk space
- Three synthetic network adapters

When the VSM virtual machine is deployed using the installer application or the VSM virtual machine template for Microsoft SCVMM, the requirements described in the following sections are correctly enforced.

VSM Interfaces

The VSM is a virtual machine that requires three virtual network interface cards (vNICs). Each vNIC has a specific function, and all are fundamental to the operation of the Cisco Nexus 1000V. To define the VSM virtual machine properties, the vNICs require the synthetic network adapter.

The VSM can be installed using either the installer application or a manual installation process. The installer application and the VSM template (provided for manual installation) help ensure that the VSM virtual machine is configured with a synthetic network adapter.

Control Interface

The control interface is primarily used for VSM high-availability communication between the primary VSM and the secondary VSM when high-availability mode is used. This interface handles low-level control packets such as heartbeats. Because of the nature of the traffic carried over the control interface, this interface is of most importance in Cisco Nexus 1000V Switch.

Some customers like to keep network management traffic in a network separate from the host management network. By default, the Cisco Nexus 1000V uses the management interface on the VSM to communicate with the VEM. However, this communication can be moved to the control interface by configuring server virtualization switch (SVS) mode to use the control interface. Please refer to the Cisco Nexus 1000V configuration guide for additional information about how to use the control interface for VSM-to-VEM communication.

Management Interface

The management interface appears as the mgmt0 port on a Cisco switch. As with the management interfaces of other Cisco switches, an IP address is assigned to mgmt0. Because Layer 2 communication is not supported between the VSM and the VEM on the Microsoft Hyper-V host as in VMware ESX, the management interface is used for all VSM-to-VEM communication by default.

The VSM administrator can use the **show interface mac address** command to see the MAC addresses assigned to control and management interfaces (Figure 16).

Figure 16. Viewing the MAC Addresses of the Control and Management Interfaces

```
Nexus1KV(config)#
Nexus1KV(config)#
Nexus1KV(config)# sh interface mac-address

-----
Interface                Mac-Address      Burn-in Mac-Address
-----
mgmt0                     001d.d8b7.1c1f   001d.d8b7.1c1f
control0                  001d.d8b7.1c1e   001d.d8b7.1c1e
Nexus1KV(config)#
Nexus1KV(config)#
```

Packet Interface

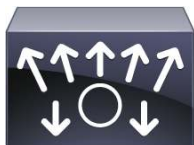
The packet interface is a traditional interface on the Cisco Nexus 1000V VSM for Microsoft Hyper-V.

Virtual Ethernet Module

The VEM provides the Cisco Nexus 1000V with network connectivity and forwarding capabilities much like a line card in a modular switching platform. Unlike multiple line cards in a single chassis, each VEM acts as an independent switch from a forwarding perspective.

The VEM is tightly integrated with the Microsoft Hyper-V hypervisor. The VEM is installed as a forwarding extension to the Microsoft Hyper-V extensible switch that runs in the Microsoft Windows server kernel (Figure 17).

Figure 17. VEM Representation



Unlike with the VSM, the VEM's resources are unmanaged and dynamic. Although the storage footprint of the VEM is fixed (approximately 6.4 MB of disk space), RAM use on the Microsoft Hyper-V host is variable, based on the configuration and scale of the Cisco Nexus 1000V deployment. In a typical configuration, each VEM can be expected to require 10 to 50 MB of RAM, with an upper limit of 150 MB for a fully scaled solution with all features turned on and used to their design limits.

Each instance of the Cisco Nexus 1000V is typically composed of two VSMs (in a high-availability pair) and one or more VEMs. The maximum number of VEMs supported by a VSM is 64.

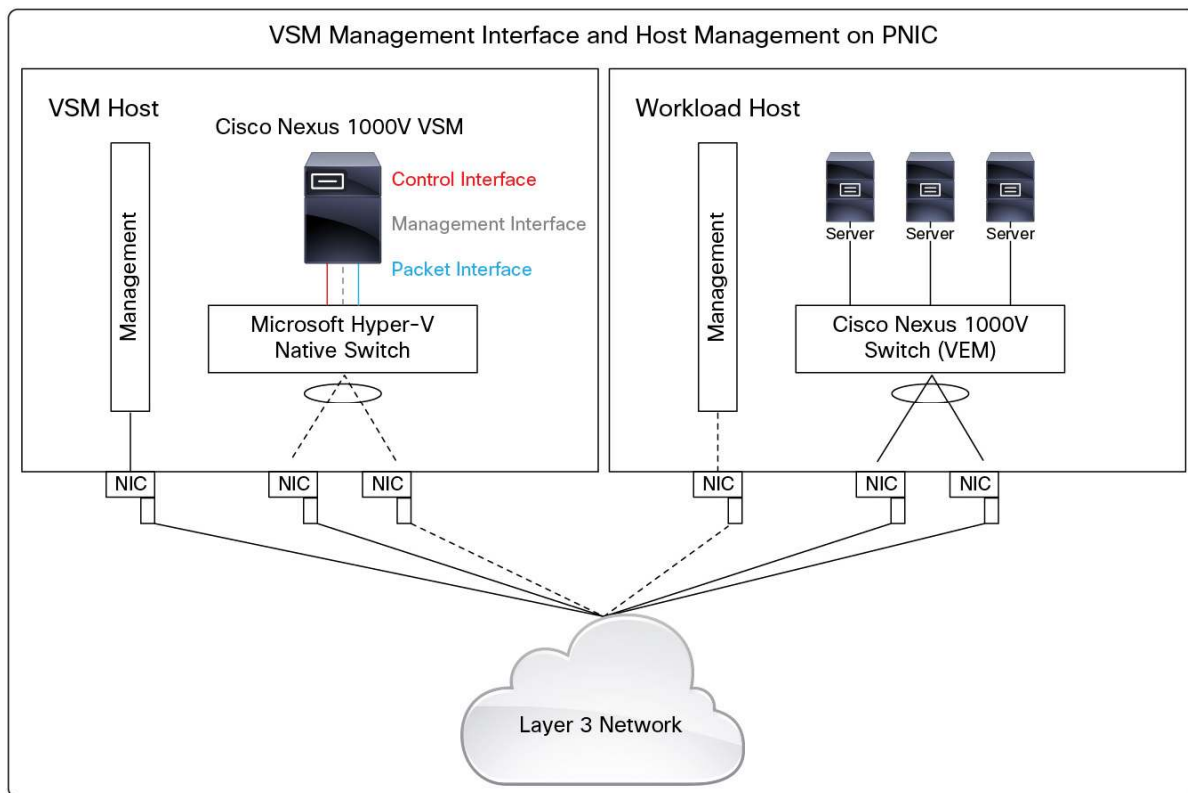
VSM-to-VEM Communication

The Cisco Nexus 1000V VSM needs to communicate with the VEM to program and monitor virtual switch ports on the Logical Switch instances (VEMs). The Cisco Nexus 1000V Switch for Microsoft Hyper-V supports only Layer 3 communication between the VSM and the VEMs. In Layer 3 mode, the control information is exchanged as Layer 3 IP packets between the VSM and VEM. Because the VSM and VEM exchange IP traffic, the VSM and VEMs can be on different subnets as long as routing between the subnets is configured correctly. Layer 3 mode deployment leads to greater deployment flexibility.

Note: Unlike the Cisco Nexus 1000V Switch for VMware ESX, which supported both Layer 3 and Layer 2 modes, the Cisco Nexus 1000V Switch for Microsoft Hyper-V supports only Layer 3 mode.

By default the management interface on the VSM (mgmt0 interface) is used for VSM-to-VEM communication. In Figure 18, the dotted lines show the path taken by VSM-to-VEM communication.

Figure 18. Management Interface on the VSM and the Microsoft Hyper-V Host Management Physical NIC Used for Connectivity



Some administrators prefer to keep management and control traffic on separate networks on the VSM. This separation can be achieved by using the control0 interface on the VSM to communicate with the physical management interface on the Microsoft Hyper-V hosts. It can be configured using the following **svs** command:

```
Nexus1000v(config-svs-domain)# svs mode L3 interface[Mgmt0|Control0]
```

Following is an example showing how to configure the Control0 interface for communication:

```
Nexus1000v(config)# int control 0
Nexus1000v(config-if)# ip address 192.168.150.10 255.255.255.0

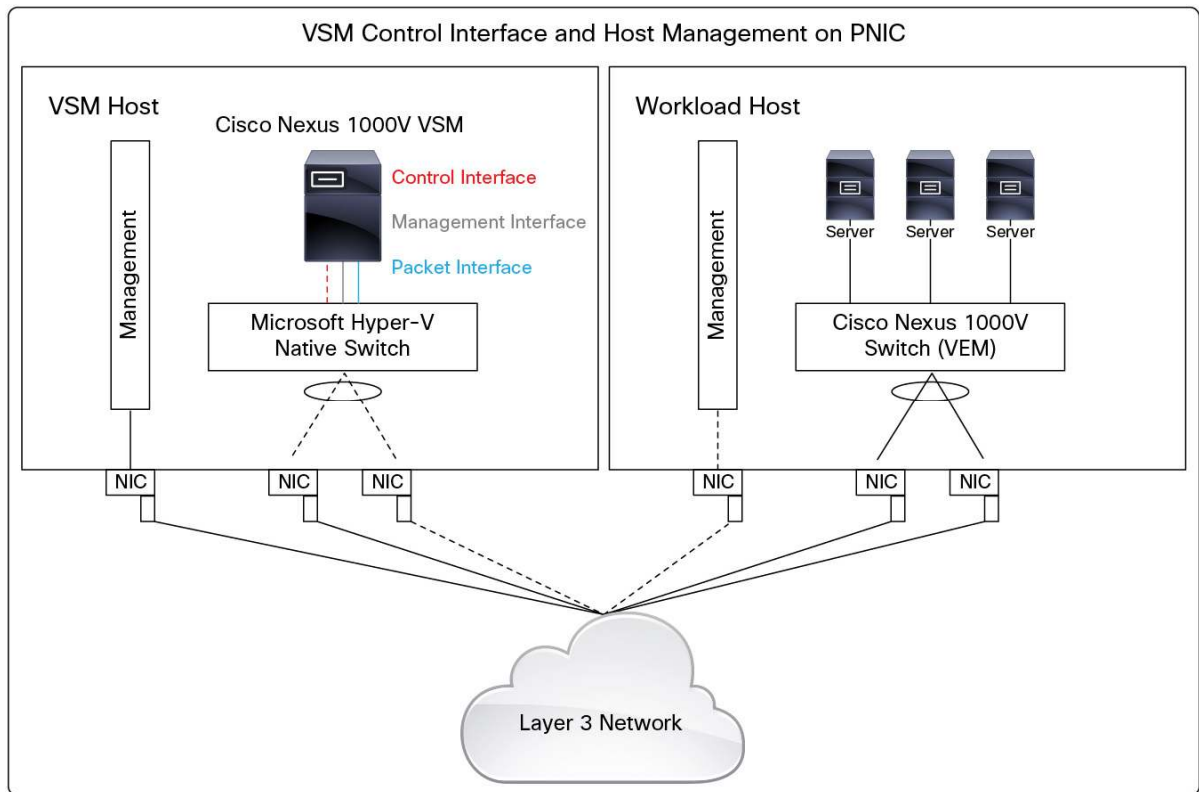
Nexus1000v(config-if)# svs-domain
Nexus1000v(config-svs-domain)# svs mode L3 interface control0
```

After entering this code, try to ping the 192.168.150.10 interface from the Microsoft Hyper-V hosts. You should not be able to because a default route has not yet been set for the default VRF instance. The following CLI shows how to set the default router

```
Nexus1000v(config)# vrf context default
Nexus1000v(config)# ip route 0.0.0.0/0 192.168.150.1
```

Figure 19 shows the control0 interface used to communicate with the VEM.

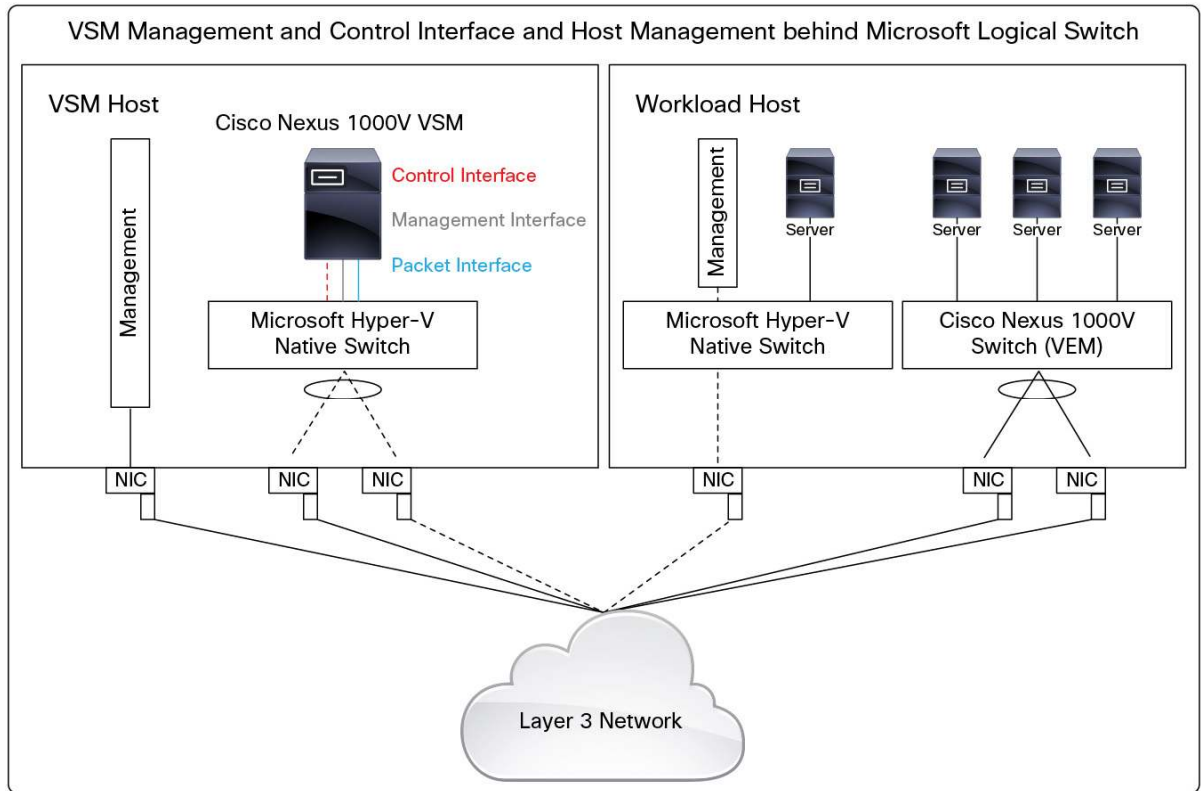
Figure 19. Control0 Interface on the VSM and the Microsoft Hyper-V Host Management Physical NIC Used for Connectivity



Note: The Cisco Nexus 1000V relies on proxy Address Resolution Protocol (ARP) when control0 interface is used for VEM-to-VSM communication. Proxy ARP must be enabled on the control0 VLAN in the gateway router.

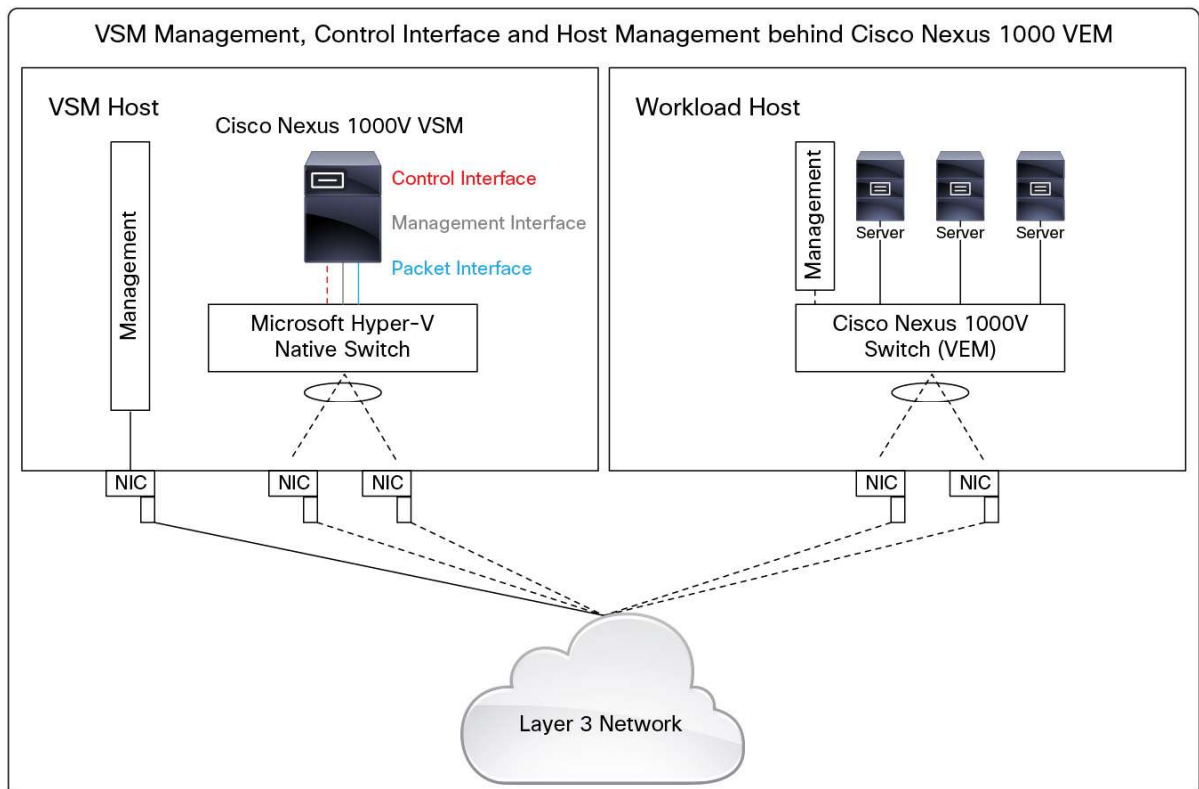
Some customers prefer to move the Microsoft Hyper-V host management interface behind a Microsoft virtual switch and share the physical interface with other virtual machines. In this scenario, no special Cisco Nexus 1000V configuration is needed to enable VSM-to-VEM communication (Figure 20).

Figure 20. Management and Control0 Interface on the VSM and the Virtual Management NIC Connected to the Microsoft Logical Switch



In some scenarios, customers may want to share the physical NICs (pNICs) on the host between management and workload virtual machines (storage, live migration, and cluster heartbeat traffic is also usually shared in these cases) when limited pNICs are available per Microsoft Hyper-V host. In such scenarios, management can be moved behind the Cisco Nexus 1000V (VEM), as shown in Figure 21.

Figure 21. Management and Control0 Interface on the VSM and the Virtual Management NIC Behind the Cisco Nexus 1000V



Microsoft Hyper-V has a special requirement for management vNICs. A management vNIC can be created only during virtual switch (native switch or Cisco Nexus 1000V switch) creation on a Microsoft Hyper-V host. In other words, a management vNIC cannot be created and deployed to an already existing virtual switch (vSwitch). This management vNIC creation activity must be performed during vSwitch creation.

To create the topology shown here on a Microsoft Hyper-V host, the Microsoft SCVMM administrator creates a Cisco Nexus 1000V Logical Switch on the Microsoft Hyper-V host with the pNIC currently used for management connectivity. As part of this process, the administrator must also create a management vNIC and apply a system port profile to the vNIC. When the management vNIC is moved behind the Cisco Nexus 1000V (VEM), the network segment (VLAN) and the port profile used by the management vNIC must be marked as “system network segment” and “system port profile,” respectively. On the physical NIC, the uplink profile with the system VLAN configuration must be applied. The Microsoft Hyper-V host automatically copies the IP address and Domain Name System (DNS) setting from the pNIC to the vNIC that is created.

Note: Typically, a pNIC does not use VLAN tags for communication; therefore, while moving the management vNIC behind the Cisco Nexus 1000V, set the management VLAN on the uplink profile to native. Failure to do so may lead to loss of connectivity to the host.

Note: It is highly recommended that user adds only the management pNIC to the Cisco Nexus 1000V while moving the management NIC behind the VEM. Other pNICs can be added to the Cisco Nexus 1000V after the module successfully attaches to the VSM

1. Mark the virtual Ethernet port profile as a system port profile:

```
switch# configure terminal
switch(config)# port-profile type vethernet AccessProfile
switch(config-port-prof)# no shutdown
switch(config-port-prof)# state enabled
switch(config-port-prof)# system port-profile
switch(config-port-prof)# publish port-profile
switch(config-port-prof)# end
```

2. Mark the network segment as a system network segment.

```
switch(config)# nsm network segment VMNetwork1
switch(config-net-seg)# switchport access vlan 100
switch(config-net-seg)# member-of network segment pool InternalSF
switch(config-net-seg)# system network segment
switch(config-net-seg)# publish network segment
switch(config-net-seg)# exit
```

3. Mark the uplink network as a system uplink network.

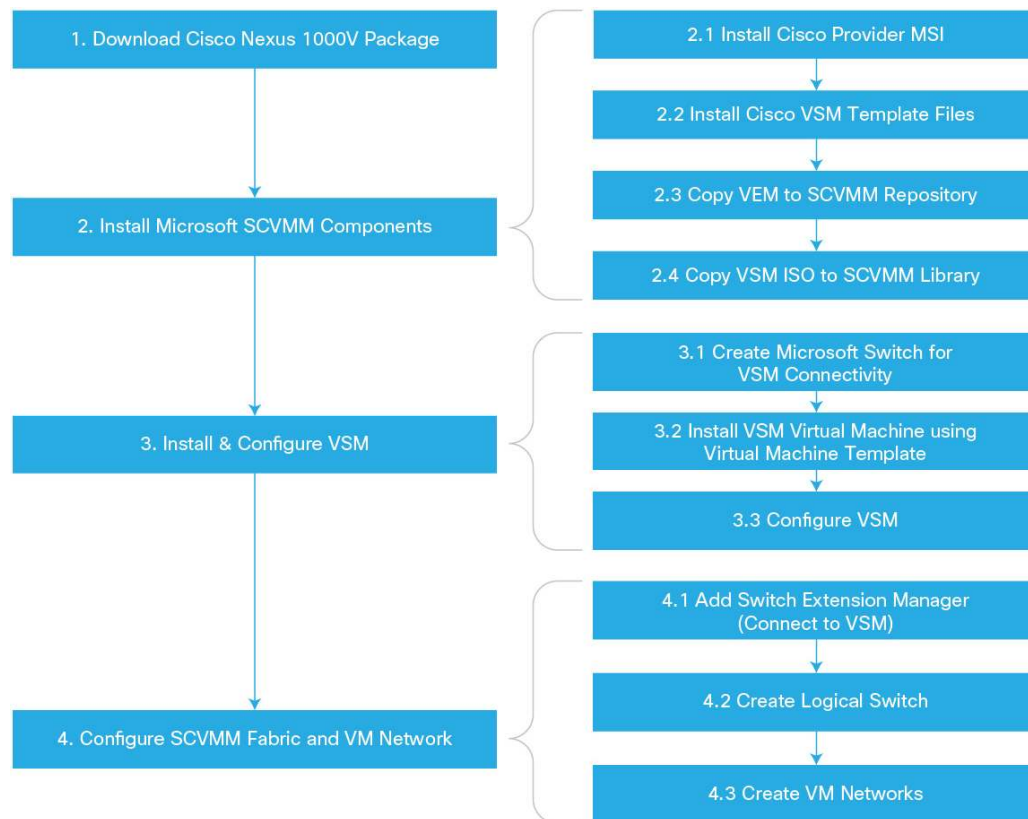
```
switch(config)# nsm network uplink Channel
switch(config-uplink-net)# allow network segment pool InternalSF
switch(config-uplink-net)# system network uplink
switch(config-uplink-net)# publish network uplink
```

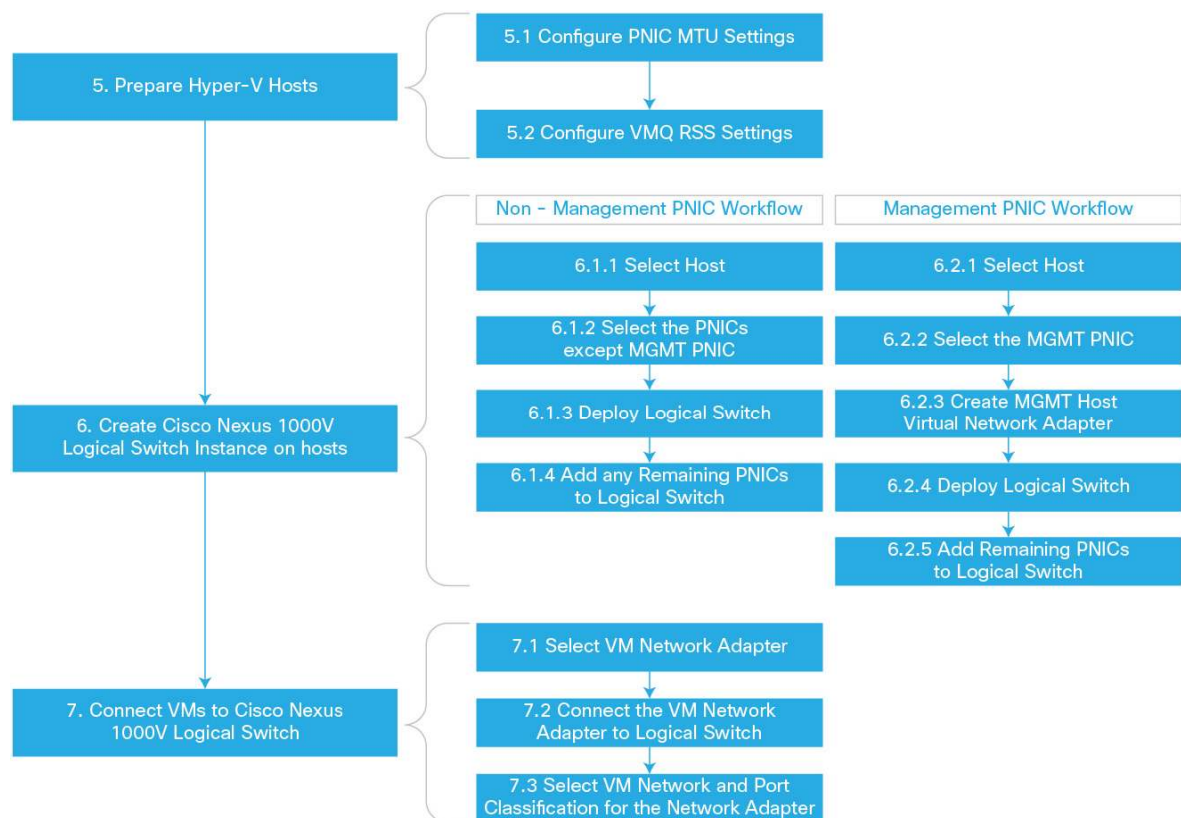
Note: The topology shown here, moving the management NIC behind the Cisco Nexus 1000V, is a supported configuration. However, this is not a recommended topology for this release. This recommendation is consistent with Microsoft's recommendation that management, storage, live migration, and cluster traffic go directly to the pNIC whenever possible.

Cisco Nexus 1000V Switch Installation

Installation of the Cisco Nexus 1000V Switch is beyond the scope of this document. Figure 22 shows the Cisco Nexus 1000V installation steps at a high level for conceptual completeness. For guidance and detailed instructions about installation, please refer to the Cisco Nexus 1000V installation guide.

Figure 22. Installation Steps for Cisco Nexus 1000V Switch for Microsoft Hyper-V





Cisco Nexus 1000V Switch Features

Switch Forwarding

In many ways, the Cisco Nexus 1000V Switches are similar to physical Ethernet switches. For packet forwarding, the Cisco Nexus 1000V Switch uses the same techniques that other Ethernet switches apply, with a MAC address-to-port mapping table used to determine the location to which packets should be forwarded.

The Cisco Nexus 1000V Switch maintains forwarding tables in a slightly different manner than other modular switches. Unlike physical switches with a centralized forwarding engine, each VEM maintains a separate forwarding table. There is no synchronization between forwarding tables on different VEMs. In addition, there is no concept of forwarding from a port on one VEM to a port on another VEM. Packets destined for a device not local to a VEM are forwarded to the external network, which in turn may forward the packets to a different VEM.

MAC Address Learning

This distributed forwarding model in a centrally managed switch is demonstrated by the way the Cisco Nexus 1000V handles MAC address learning. A MAC address can be learned multiple times within a single Cisco Nexus 1000V Switch in either of two ways: statically or dynamically. Static entries are automatically generated for virtual machines running on the VEM; these entries do not time out. For devices not running on the VEM, the VEM can learn a MAC address dynamically, through the pNICs in the server.

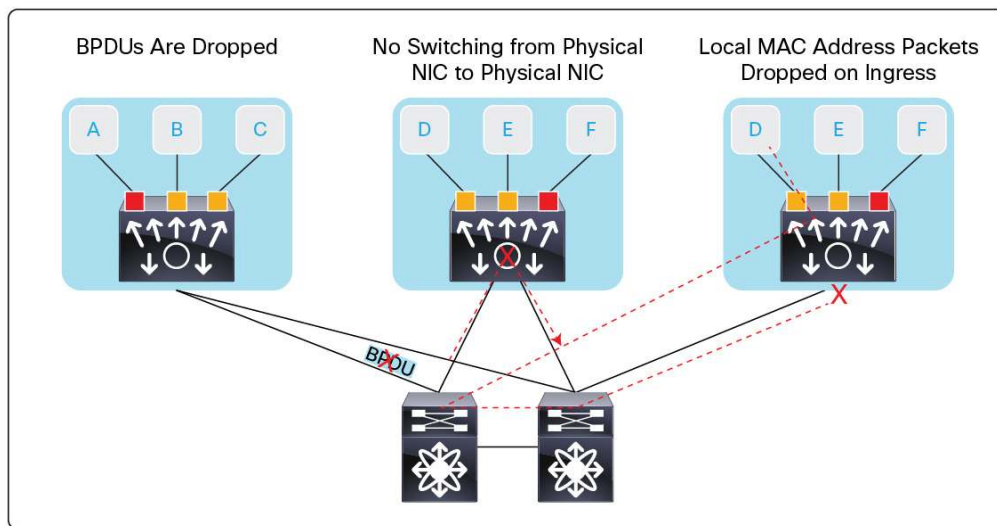
Each VEM maintains a separate MAC address table. Thus, a single Cisco Nexus 1000V Switch may learn a given MAC address multiple times: as often as once per VEM. For example, one VEM may be hosting a virtual machine, and the virtual machine's MAC address will be statically learned on the VEM. A second VEM, in the same Cisco

Nexus 1000V Switch, may learn the virtual machine's MAC address dynamically. Thus, within the Cisco NX-OS CLI, you may see the virtual machine's MAC address twice: a dynamic entry and a static entry.

Loop Prevention

Another differentiating characteristic of the Cisco Nexus 1000V is that it does not run Spanning Tree Protocol. Although this may seem to be a significant departure from other Ethernet switches, potentially causing catastrophic network loops, in reality the Cisco Nexus 1000V implements a simple and effective loop-prevention strategy that does not require Spanning Tree Protocol (Figure 23).

Figure 23. Built-in Loop Prevention Capabilities



Because the Cisco Nexus 1000V does not participate in Spanning Tree Protocol, it does not respond to Bridge Protocol Data Unit (BPDU) packets, nor does it generate them. BPDU packets that are received by Cisco Nexus 1000V Switches are dropped.

The Cisco Nexus 1000V uses a simple technique to prevent loops. Like a physical Ethernet switch, the Cisco Nexus 1000V Switch performs source and destination MAC address lookups to make forwarding decisions. The VEM applies loop-prevention logic to every incoming packet on Ethernet interfaces. This logic is used to identify potential loops. Every ingress packet on a physical Ethernet interface is inspected to help ensure that the destination MAC address is internal to the VEM. If the source MAC address is internal to the VEM, the Cisco Nexus 1000V Switch will drop the packet. If the destination MAC address is external, the switch will drop the packet, preventing a loop back to the physical network.

Note: The Cisco Nexus 1000V prevents loops between the VEMs and the first-hop access switches without the use of Spanning Tree Protocol. However, this feature does not mean that Spanning Tree Protocol should be disabled on any access switches. Spanning Tree Protocol is still required by access switches to prevent loops elsewhere in the physical topology.

Spanning Tree Protocol goes through a series of states on each interface as it tries to build the network tree. This process causes downtime on each interface when Spanning Tree Protocol needs to converge. This process is unnecessary for ports connected to Cisco Nexus 1000V Switches. By using the PortFast feature on a switch port, a Cisco switch can suppress the progression of Spanning Tree Protocol states and move straight to a

forwarding state. PortFast is configured per interface and should be enabled on interfaces connected to a VEM, along with BPDU guard and BPDU filtering. Filtering BPDUs at the physical switch port will enhance VEM performance by avoiding unnecessary processing at the VEM uplink interfaces.

Switch Port Interfaces

The Cisco Nexus 1000V supports multiple switch-port types for internal and external connectivity: virtual Ethernet (vEth), Ethernet (Eth), and PortChannel (Po). The most common port type in a Cisco Nexus 1000V environment is the vEth interface, which is a new concept. This interface type represents the switch port connected to a virtual machine's vNIC or management or live migration vNIC.

A vEth interface has several characteristics that differentiate it from other interface types. Besides the obvious fact that vEth interfaces are virtual and therefore have no associated physical components, the interface naming convention is unique. Unlike a traditional Cisco interface, a vEth interface's name does not indicate the module with which the port is associated. Whereas a traditional physical switch port may be notated as GigX/Y, where X is the module number and Y is the port number on the module, a vEth interface is notated like this: vEthY. This unique notation is designed to work transparently with Live Migration, keeping the interface name the same regardless of the location of the associated virtual machine.

The second characteristic that makes a vEth interface unique is its transient nature. A given vEth interface appears or disappears based on the status of the virtual machine connected to it. The mapping of a virtual machine's vNIC to a vEth interface is static. When a new virtual machine is created, a vEth interface is also created for each of the virtual machine's vNICs. The vEth interfaces will persist as long as the virtual machine exists. If the virtual machine is temporarily down (the guest OS is shut down), the vEth interfaces will remain inactive but still bound to that specific virtual machine. If the virtual machine is deleted, the vEth interfaces will become available for connection to newly provisioned virtual machines.

The Cisco Nexus 1000V contains two interface types related to the Physical Ethernet interfaces (pNICs) on a Microsoft Hyper-V host. An Ethernet, or Eth, interface is represented in standard Cisco interface notation (EthX/Y) using the Cisco NX-OS naming convention "Eth" rather than a speed such as "Gig" or "Fast," as is the custom in Cisco IOS Software. These Eth interfaces are module specific and are designed to be fairly static within the environment.

PortChannels are the third interface type supported by the Cisco Nexus 1000V. A PortChannel is an aggregation of multiple Eth interfaces on the same VEM.

Note: PortChannels are not created by default and must be explicitly defined.

Opaque Data

Opaque data is a collection of Cisco Nexus 1000V configuration parameters maintained by the VSM and Microsoft SCVMM server when the link between the two is established. The opaque data contains configuration details that each VEM needs to establish connectivity to the VSM during VEM installation.

Among other content, the opaque data contains:

- Switch domain ID
- Switch name
- Control and packet VLAN IDs
- System port profiles

When a new VEM is online, either after initial installation or upon restart of a Microsoft Hyper-V host, it is an unprogrammed line card. To be correctly configured, the VEM needs to communicate with the VSM. The Microsoft SCVMM server automatically sends the opaque data to the VEM, which the VEM uses to establish communication with the VSM and download the appropriate configuration data.

Port Profiles

Port profiles are network policy templates that enable the VSM administrator to define network policy once and then apply it all interfaces requiring identical configuration. Any change made to the port profile is automatically applied to all interfaces that are using the port profile.

The Cisco Nexus 1000V Switch supports creation of two types of port profiles: vEth and Eth port profiles. Virtual Ethernet port profiles are applied to vEth interfaces, and Ethernet port profiles are applied to server uplink Eth interfaces.

The Cisco Nexus 1000V Switch for Microsoft Hyper-V also incorporates the new concept of dynamic port profiles, which is explained in the following sections.

Virtual Ethernet Port Profiles

Virtual Ethernet port profiles are the primary mechanisms by which network policy is defined and applied to the Microsoft Hyper-V switch's vEth interfaces.

A port profile can be applied on a virtual interface using the **vethernet** keyword for the port profile type or on a physical interface using the **ethernet** keyword for the port profile type. If no keyword is specified, the vEth type (**vethernet**) is used by default.

A port profile is a collection of interface-level configuration commands that are combined to create a complete network policy.

The port profile concept is new, but the configurations in port profiles use the same Cisco syntax that is used to manage switch ports on traditional switches. The VSM administrator:

1. Defines a new port profile in switch configuration mode
2. Applies Interface-level configuration commands
3. Enables the port profile
4. Publishes the port profile to Microsoft SCVMM

The process of publishing the port profile enables the VSEM defined for Microsoft SCVMM to pull the port profile configuration into Microsoft SCVMM. The Microsoft SCVMM administrator can then create a port classification that uses the port profile.

Note: Publishing the port profile does not automatically make it available to Microsoft SCVMM. The Microsoft SCVMM administrator must either refresh the switch extension or wait for Microsoft SCVMM to autorefresh it (refresh occurs every 30 minutes). This behavior is different from that for Cisco Nexus 1000V and VMware vCenter, in which creation of a port profile results in instantaneous creation of a port group.

A vEth profile is a port profile that can be applied on virtual machines and on the Microsoft Hyper-V host's virtual interfaces such as management, live migration, and storage vNICs.

Figure 24 shows a vEth port profile; in this example, an IP access list is applied to a port profile.

Figure 24. vEth Port Profile Example

```
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# ip access-list Restrict-ContractorVM
Nexus1000V(config-acl)# 10 deny ip 20.20.20.22/32 20.20.20.20/32
Nexus1000V(config-acl)# 20 permit ip any any
Nexus1000V(config-acl)# exit
Nexus1000V(config)#

Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# port-profile type vethernet RestrictedProfile
Nexus1000V(config-port-prof)# ip port access-group Restrict-ContractorVM in
Nexus1000V(config-port-prof)# no shutdown
Nexus1000V(config-port-prof)# state enabled
Nexus1000V(config-port-prof)# publish port-profile
Nexus1000V(config-port-prof)# exit
Nexus1000V(config)#
Nexus1000V(config)#
```

Ethernet Port Profiles

Port profiles are not used only to manage vEth configuration; they are also used to manage the pNICs in a Microsoft Hyper-V host. When a port profile is defined, the VSM administrator determines whether the profile will be used to manage vEth interfaces or pNICs. By default, the port profile is assumed to be used for vEth management.

To define a port profile for use on pNICs, the network administrator applies the **ethernet** keyword to the profile. When this option is used, the port profile is available only to the server administrator to apply to pNICs in a Microsoft Hyper-V server.

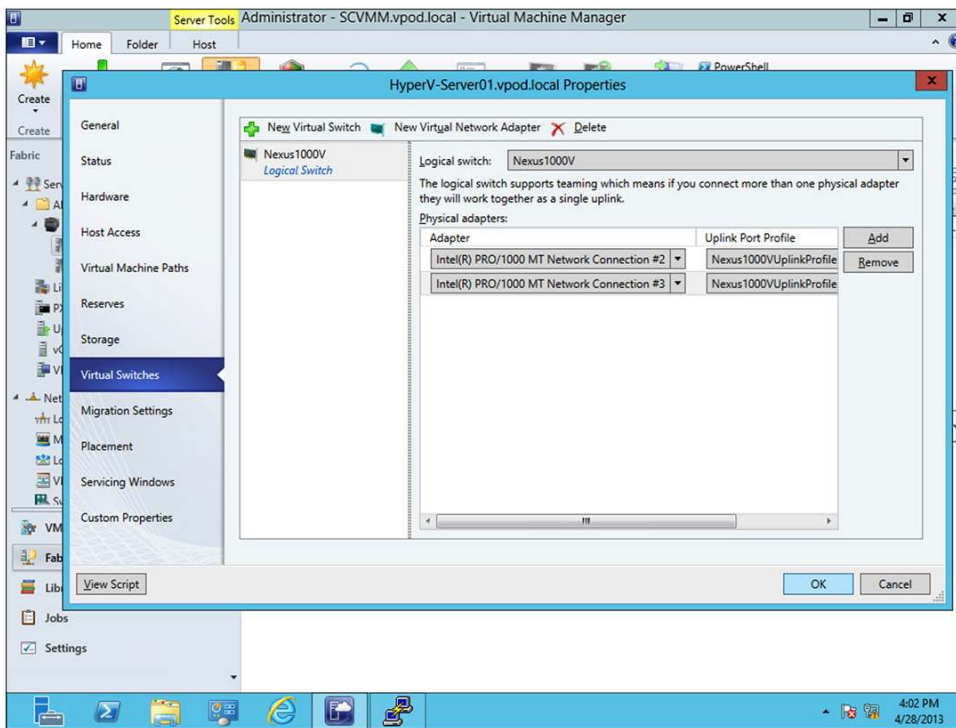
Figure 25 shows an Eth port profile using virtual PortChannel (vPC) host mode. For more information about configuring port profiles, please refer to the Cisco Nexus 1000V configuration guide.

Figure 25. Eth Port Profile Example

```
Nexus1KV(config)#
Nexus1KV(config)# port-profile type ethernet UplinkProfile
Nexus1KV(config-port-prof)# channel-group auto mode on mac-pinning
Nexus1KV(config-port-prof)# no shutdown
Nexus1KV(config-port-prof)# state enabled
Nexus1KV(config-port-prof)# exit
Nexus1KV(config)#
Nexus1KV(config)#
```

Uplink port profiles are applied to a pNIC when a Microsoft Hyper-V host is first added to the Cisco Nexus 1000V Switch. The Microsoft SCVMM administrator is presented with a dialog box in which the administrator selects the pNICs to be associated with the VEM and the specific uplink port profiles to be associated with the pNICs. In addition, the Microsoft SCVMM administrator can apply uplink port profiles to interfaces that are added to the VEM after the host has been added to the switch (Figure 26).

Figure 26. Applying an Uplink Port Profile to Server Uplink Interfaces During Logical Switch Creation



Network Segments

The **network segment** command is a new command introduced in the Cisco Nexus 1000V Switch for Microsoft Hyper-V. Network segments are used to create Layer 2 networks on the VSM. In the first release of the Cisco Nexus 1000V Switch for Microsoft Hyper-V, only VLAN-based network segments are supported. Other segmentation technology is not supported in this release.

The Microsoft SCVMM networking model associates each network segment with a network site. When the network segment is created on the VSM, the VSM administrator must also specify the network segment pool (network site) to which the segment belongs. Figure 27 shows a network segment created using Vlan10. This network segment can then be used by the Microsoft SCVMM administrator to create a DMZ virtual machine network.

Figure 27. Creating a Network Segment

```
Nexus1000V(config)#
Nexus1000V(config)# nsm network segment Vlan10
Nexus1000V(config-net-seg)# member-of network segment pool DM2-SFO
Nexus1000V(config-net-seg)# switchport access vlan 10
Nexus1000V(config-net-seg)# switchport mode access
Nexus1000V(config-net-seg)# ip pool import template Vlan10IPPool
Nexus1000V(config-net-seg)# publish network segment
Nexus1000V(config-net-seg)# exit
Nexus1000V(config)#
```

Note: The Cisco Nexus 1000V Switch for Microsoft Hyper-V does not support VLAN creation using the **vlan <id>** command. The VSM administrator must use the **nsm network segment** command.

Note: In the first release of the Cisco Nexus 1000V Switch for Microsoft Hyper-V, trunk vEth interfaces are not supported. Hence, the network segments switch-port mode must always be set to “access.”

Dynamic Port Profiles

The Cisco Nexus 1000V introduces the concept of the dynamic port profile for Microsoft Hyper-V to support policy and network separation design in Microsoft SCVMM. When the Microsoft SCVMM administrator deploys a virtual machine to a Cisco Nexus 1000V port classification and virtual machine network, a dynamic port profile is created on the Cisco Nexus 1000V VSM and is applied to the virtual switch port on which the virtual machine is deployed. These dynamic port profiles are shared by all virtual machines that have the same virtual machine network and port classification. These port profiles are internal to the Cisco Nexus 1000V VSM and must not be edited or deleted by the Cisco Nexus 1000 VSM administrator.

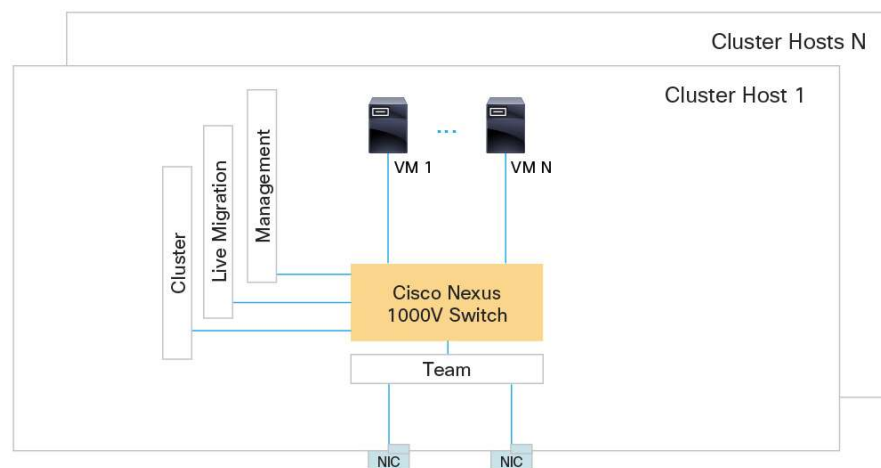
System Network Segments

System network segments can be used to put virtual switch ports on the Cisco Nexus 1000V VEM in perpetual forwarding mode. When system network segments are not used, virtual switch ports on the VEM must be programmed by the VSM before they will start forwarding traffic.

In certain deployments, the administrator may want to always allow traffic on certain switch ports. Figure 28 shows an of such a deployment. In this topology, the management vNIC of the Microsoft Hyper-V host is attached to a Cisco Nexus 1000V virtual switch port. When the Microsoft Hyper-V host restarts, all the switch ports are disabled for forwarding. In this scenario, the Cisco Nexus 1000V VSM will attempt to reprogram all the switch ports on the Microsoft Hyper-V host. However, the VSM will fail because the Microsoft Hyper-V host will not be reachable because it is not in the forwarding state.

In such a scenario, the administrator should define the network segment that is used to manage the Microsoft Hyper-V hosts as a system network segment. Because all ports in a system network segment are in perpetual forwarding mode, the VSM will be able to reprogram the VEMs across reboot and VSM-to-VEM reconnections.

Figure 28. Management Network Adapter Behind Cisco Nexus 1000V Switch



To enable perpetual forwarding on switch ports, the network segment and the uplink network must be configured, respectively, as “system network segment” and “system uplink network,” as shown in the following example:

```
switch(config)# nsm logical network Internal
switch(config-logical-net)# description network for host/VMs behind the DMZ
switchM(config-logical-net)# exit

switch(config)# nsm network segment pool InternalSF
```



```

switch(config-net-seg-pool)# member-of logical network Internal
switch(config-net-seg-pool)# exit

switch(config)# nsm network segment VMNetwork1
switch(config-net-seg)# switchport access vlan 100
switch(config-net-seg)# member-of network segment pool InternalSF
switch(config-net-seg)# system network segment
switch(config-net-seg)# publish network segment
switch(config-net-seg)# exit

switch(config)# nsm network uplink Channel
switch(config-uplink-net)# allow network segment pool InternalSF
switch(config-uplink-net)# system network uplink
switch(config-uplink-net)# publish network uplink

```

Policy Mobility

Network policies enforced by a port profile follow the virtual machine throughout its lifecycle, whether the virtual machine is being migrated from one server to another, suspended, hibernated, or restarted. In addition to migrating the policy, the Cisco Nexus 1000V Switches move the virtual machine's network state, such as the port counters and flow statistics. Virtual machines participating in traffic monitoring activities, such as Cisco NetFlow or Encapsulated Remote Switched Port Analyzer (ERSPAN), can continue these activities uninterrupted by Microsoft live migration operations.

Licensing

The Cisco Nexus 1000V Switch for Microsoft Hyper-V is shipped in two editions: Essential and Advanced. A new CLI command **show switch edition** is provided to display the current switch edition and other licensing information.

In the two-tier licensing model, the software image is the same for both editions. You can switch between the Essential edition and the Advanced edition at any time. The switch edition configuration is global. The entire switch (the supervisor and all modules) uses either in the Essential edition or the Advanced edition.

In the two-tier licensing approach, the licenses are checked out only if the switch edition is Advanced. In the Essential edition, the license checkout process is skipped. The modules automatically transition to the licensed state.

The following features are available as advanced features that require licenses: Cisco TrustSec[®] capability, Dynamic Host Configuration Protocol (DHCP) snooping, IP source guard, and Dynamic ARP Inspection (DAI). The DHCP snooping, IP source guard, and DAI features can be enabled using the **feature** command:

```

nexus1000v(config)# svs switch edition < essential | advanced >
nexus1000v# show switch edition

```

Please refer to the Cisco Nexus 1000V licensing guide for additional information.

Cisco Nexus 1000V Switch Network Design

Deployment Topology 1: Recommended

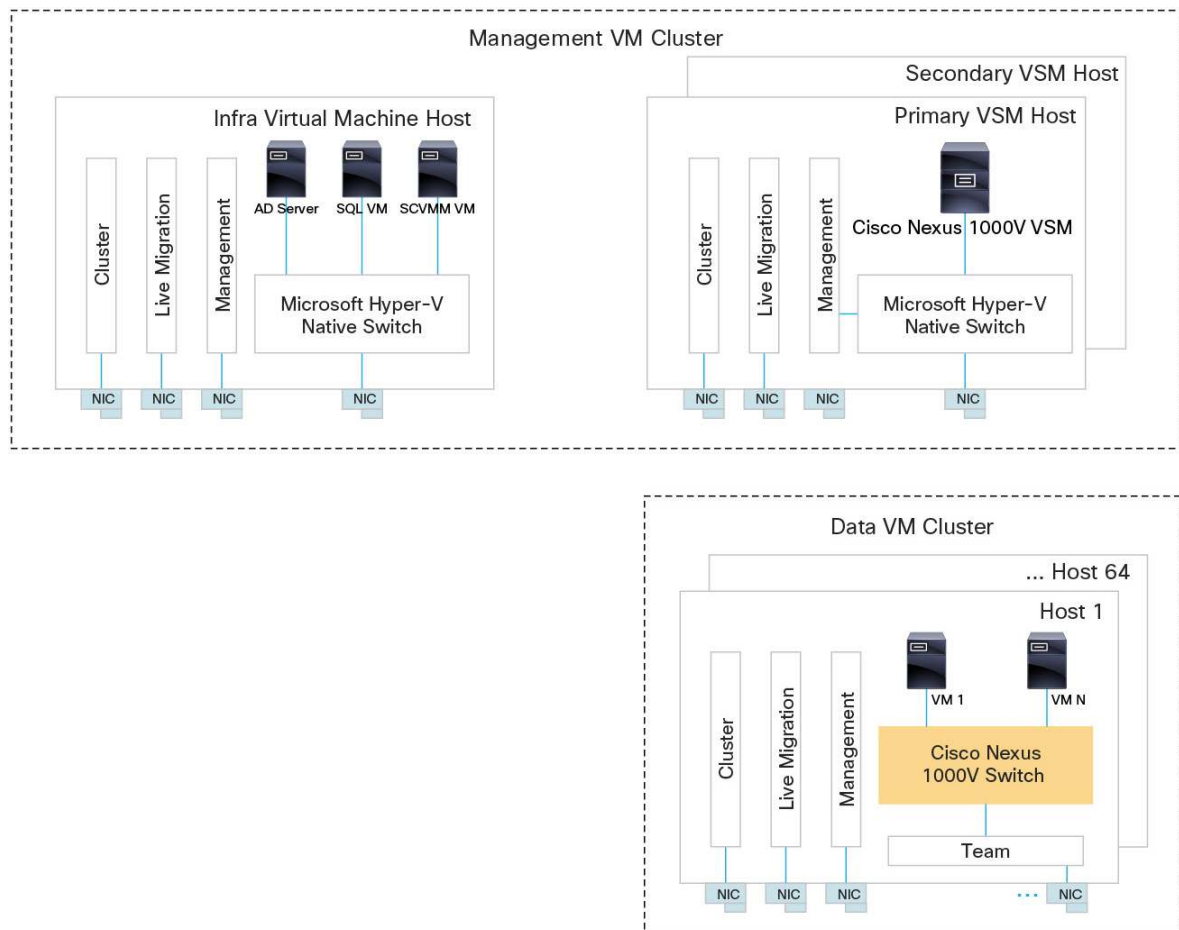
Figure 29 shows the recommended topology for Cisco Nexus 1000V Switch for Microsoft Hyper-V deployment. This topology is used for larger Microsoft Hyper-V deployments in which infrastructure virtual machines are

separated from workload virtual machines (running line-of-business [LoB] applications) by running them on separate Microsoft Hyper-V hosts or clusters.

Management Microsoft Hyper-V Clusters and Hosts

In this Microsoft Hyper-V deployment topology, infrastructure virtual machines are deployed on a separate set of hosts or on an infrastructure Microsoft Windows failover cluster. Examples of infrastructure virtual machines are Microsoft Active Directory servers, DNS servers, SQL servers, and Microsoft SCVMM and other Microsoft System Center roles. The Cisco Nexus 1000V VSM virtual machine should also be deployed on an infrastructure host or cluster. The Cisco Nexus 1000V Logical Switch (VEM) is not created on the infrastructure hosts; instead, the native Microsoft Hyper-V switch is used.

Figure 29. Deployment Topology 1 (Recommended)



VSM High Availability

Cisco Nexus 1000V VSM deployment in high-availability mode is recommended. When the Cisco Nexus 1000V is deployed in high-availability mode, one of the virtual machine functions as the active VSM, and the other is in hot-standby mode. VSM virtual machines can be deployed on standalone Microsoft Hyper-V hosts, in Microsoft Windows failover cluster, or on Cisco Nexus 1010 and 1110 devices. For information about how to deploying the Cisco Nexus 1000V VSM on the Cisco Nexus 1010 and 1110, please refer to the Cisco Nexus 1010/1110 documentation.

You should use Cisco NX-OS high-availability mode to help ensure availability of the VSM virtual machine. When the active VSM or the Microsoft Hyper-V host that is running the active VSM goes down, the standby VSM becomes active. For additional information about VSM high availability, please refer to Cisco Nexus 1000V high-availability and redundancy documentation.

Because Cisco NX-OS high-availability mode is used to help ensure high availability of the VSM, deployment of the VSM virtual machine on standalone Microsoft Hyper-V hosts is preferred. If the VSM virtual machine is deployed in a failover cluster, then the VSM virtual machines must **not** be made highly available; otherwise, VSM failover may not occur smoothly.

Layer 3 Mode

Cisco Nexus 1000V supports on Layer 3 communication between the VSM and the VEMs (Cisco Nexus 1000V Switches for Microsoft Hyper-V on Microsoft Hyper-V hosts): that is, IP traffic is used for VSM-to-VEM communication. VSM-to-VEM communication flows between the management interface on the Cisco Nexus 1000V VSM and the Microsoft Hyper-V host management interfaces. Because the control traffic between the VSM and the VEM is based on Layer 3 IP packets, the VEM and the VSMs can be in different IP subnets. VLANs and IP routing must be configured correctly on the physical routers and switches to help ensure VSM-to-VEM communication.

Data Virtual Machine Cluster

The Cisco Nexus 1000V Logical Switch must be created only on Microsoft Hyper-V hosts that run workload virtual machines. As shown earlier in Figure 29, the Cisco Nexus 1000V Logical Switch (VEM) is not created on infrastructure hosts and is created only on workload hosts.

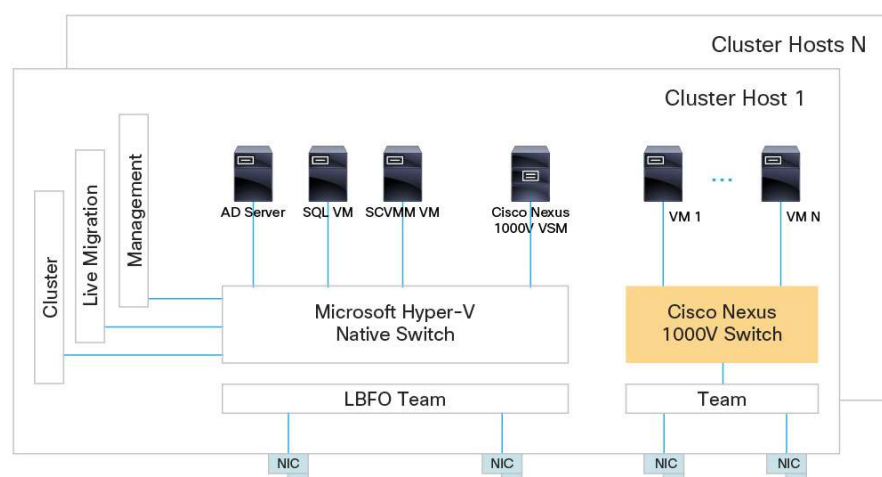
Management, Live Migration, and Cluster Traffic

Microsoft recommends that critical traffic such as host management, storage (Small Computer System Interface over IP [iSCSI]), live migration, and cluster traffic use the pNIC on a Microsoft Hyper-V host for reliability and performance. On the basis of this recommendation, use the Cisco Nexus 1000V Logical Switch only for workload virtual machines, and flowing infrastructure traffic directly to the pNIC.

Deployment Topology 2: For Customers with Limited pNICs per Microsoft Hyper-V Host

For customers who cannot implement separate infrastructure and workload clusters and who have a limited number of pNICs, the topology shown in Figure 30 is recommended.

Figure 30. Deployment Topology 2



In this topology, two virtual switches are created per Microsoft Hyper-V host: a native Microsoft Hyper-V switch and a Cisco Nexus 1000V Logical Switch.

The infrastructure and VSM virtual machines are deployed to a native Microsoft switch. To help ensure high availability, use a Microsoft Windows load-balancing and failover (LBFO) team as the server uplink. The team will consist of at least two physical adapters. In this deployment, the LBFO team is shared among the cluster, live migration, management, and other infrastructure virtual machines. Because some of this traffic can be bandwidth intensive (storage and live migration traffic can be bandwidth intensive in bursts), you may need to configure bandwidth reservation on the native Microsoft Hyper-V switch.

The workload virtual machines are deployed on the Cisco Nexus 1000V Logical Switch. The Logical Switch will have at least two adapters connected as the switch uplinks. vPC host mode (explained in detail later in this document) is the recommended configuration for the Cisco Nexus 1000V uplinks to help ensure the high availability of the workload virtual machines.

Deploying Cisco Nexus 1000V on Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS®) is increasingly deployed in data centers because of its advantages in server virtualization environments. In server virtualization environments in general, network management is becoming increasingly complex, and it is increasingly difficult to help ensure sure that network function requirements are met. The Cisco Nexus 1000V distributed virtual switch provides a way to support these network requirements and gives the networking team the visibility needed to manage the growing virtual data center.

Some Cisco UCS functions are similar to those offered by the Cisco Nexus 1000V Switches, but with a different set of applications and design scenarios. Cisco UCS offers the capability to present adapters to physical and virtual machines directly. This solution is a hardware-based Cisco Data Center Virtual Machine Fabric Extender (VM-FEX) solution, whereas the Cisco Nexus 1000V is a software-based VN-Link solution. This document does not discuss the differences between the two solutions. This section discusses how to deploy the Cisco Nexus 1000V in a Cisco UCS blade server environment, including best practices for configuring the Cisco Nexus 1000V for Cisco UCS. It also explains how some of the advanced features of both Cisco UCS and the Cisco Nexus 1000V facilitate the recommended deployment of the solution.

Cisco Virtual Interface Card

The Cisco UCS M81KR Virtual Interface Card (VIC), VIC 1240, and VIC 1280 are virtualization-optimized Fibre Channel over Ethernet (FCoE) mezzanine cards designed for use with Cisco UCS B-Series Blade Servers. A VIC is a dual-port 10 Gigabit Ethernet mezzanine card that supports up to 128 Peripheral Component Interconnect Express (PCIe) standards-compliant virtual interfaces that can be dynamically configured so that both their interface type (NIC or host bus adapter [HBA]) and identity (MAC address and worldwide name [WWN]) are established using just-in-time provisioning. In addition, the Cisco VIC supports Cisco VN-Link technology, which adds server-virtualization intelligence to the network. Each card has a pair of 10 Gigabit Ethernet connections to the Cisco UCS backplane that support the IEEE 802.1 Data Center Bridging (DCB) function to facilitate I/O unification within these adapters. On each adapter type, one of these backplane ports is connected through 10GBASE-KR to the A-side I/O module; then that connection goes to the A-side fabric interconnect. The other connection is 10GBASE-KR to the B-side I/O module; that connection goes to the B-side fabric interconnect.

The Cisco UCS 6100 Series Fabric Interconnects operate in two discrete modes with respect to flows in Cisco UCS. The first is assumed to be more common and is called end-host mode; the second is switched mode, in which the fabric interconnect acts as a normal Ethernet bridge device. Discussion of the differences between these modes is beyond the scope of this document; however, the Cisco Nexus 1000V Switches on the server blades will operate regardless of the mode of the fabric interconnects. With respect to a Microsoft environment running the Cisco Nexus 1000V Switch for Microsoft Hyper-V, the preferred solution is end-host mode to help ensure predictable traffic flows.

Service Profile Design

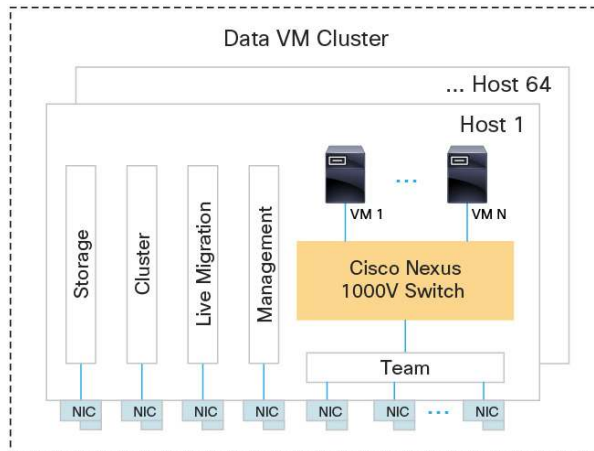
Service profiles in Cisco UCS allow the administrator to create a consistent configuration across multiple Cisco UCS server blades. The service profile definition includes server identity information such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network VLAN, physical ports, and quality-of-service (QoS) policies. For more information about how to configure service profiles, please refer to the Cisco UCS documentation.

Using service profiles, the administrator can apply security and QoS policies on multiple vEth interfaces. Virtual machines running Microsoft Windows Server 2012 on the Cisco UCS B-Series see these vNICs as physical adapters connected to the server.

When the Cisco Nexus 1000V is deployed on Cisco UCS B-Series blades, the recommended topology is the one shown in Figure 31. In this topology, the Cisco Nexus 1000V Logical Switch sees traffic only from the workload virtual machines. Because the infrastructure traffic goes directly to the vNICs, you should create separate service profiles for these two types of interfaces.

Also, because the Cisco VIC supports a minimum of 128 PCIe interfaces, you should create at least 6 vNICs per blade running Microsoft Windows Server 2012 as shown in Figure 31.

Figure 31. Cisco Nexus 1000V Switch for Microsoft Hyper-V on a Cisco UCS Blade Server



In this topology, the storage, cluster, live migration, and management traffic is sent directly to the Cisco UCS vNIC. Workload virtual machines running LoB applications are deployed on the Cisco Nexus 1000V Logical Switch. The Cisco Nexus 1000V also uses two or more vNICs as uplinks.

To support this topology, a separate service profile is created for the Cisco Nexus 1000V switch uplink vNICs. In this profile, fabric failover is disabled, as shown in Figure 32. To ensure high availability for workload virtual machines, vPC host mode is configured on the Cisco Nexus 1000V Switch uplinks. This configuration helps ensure that the uplinks are bound to a team. When a member link in the team fails, the Cisco Nexus 1000V VEM helps ensure that traffic from workload virtual machines fails over to one of the remaining links.

Figure 32. Fabric Failover Disabled for Cisco Nexus 1000V Uplink vNICs

Name: **eth0**

Use LAN Connectivity Template: ☐

MAC Address

MAC Address Assignment: **default(196/200)**

The MAC address will be automatically assigned from the selected pool.

Fabric ID: ☒ Fabric A ☐ Fabric B ☐ **Enable Failover**

VLANs

Select	Name	Native VLAN
<input checked="" type="checkbox"/>	default	<input checked="" type="radio"/>
<input checked="" type="checkbox"/>	IP-Storage-80	<input type="radio"/>
<input checked="" type="checkbox"/>	management-172	<input type="radio"/>
<input checked="" type="checkbox"/>	vm-data	<input type="radio"/>

MTU: **1500**

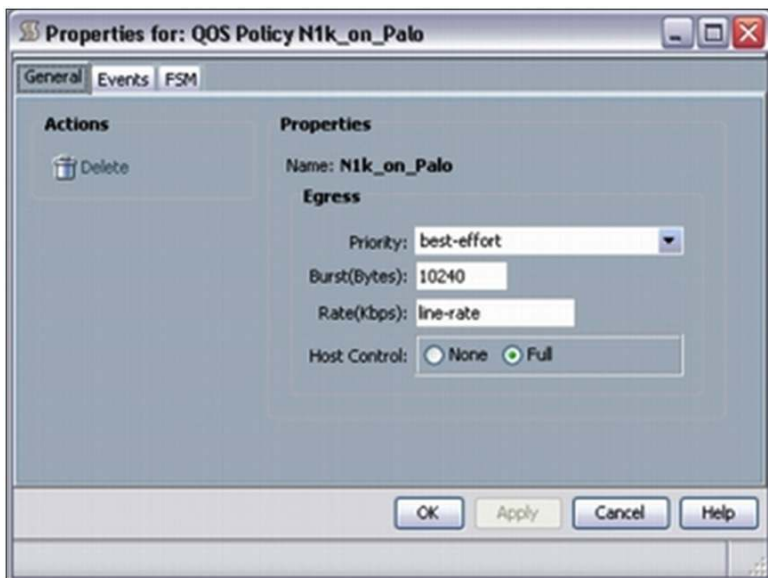
Pin Group: **<not set>**

Fabric Failover Disabled

Fabric failover is a unique capability found only in Cisco UCS that allows a server adapter to have a highly available connection to two redundant network switches (fabric interconnects) without any NIC teaming drivers or any NIC failover configuration required in the OS, hypervisor, or virtual machine. The Cisco VIC adapters - the Cisco UCS M81KR VIC, VIC 1240, and VIC 1280 adapter types - enable a fabric failover capability in which loss of connectivity on a path in use causes traffic to be remapped through a redundant path within Cisco UCS.

In the service profile definition for the vNICs used as Cisco Nexus 1000V uplinks, Host Control is set to Full (Figure 33). This setting helps ensure that class-of-service (CoS) markings from the Cisco Nexus 1000V are honored.

Figure 33. Host Control Set to Full for Cisco Nexus 1000V Uplink vNICs



In the service profile used to create infrastructure vNICs, fabric failover is enabled, and Host Control is set to None. This configuration is used because of infrastructure traffic is going directly to the Cisco UCS vNIC. For this traffic type, configure the CoS markings in the Cisco UCS service profile.

Quality of Service

Another distinguishing feature of Cisco UCS is the capability of the VIC to perform CoS-based queuing in hardware. CoS is a value marked on Ethernet packets to indicate the priority in the network. The Cisco UCS VIC has eight traffic queues, which use CoS values of 0 through 7. The VIC also allows the network administrator to specify a minimum bandwidth that must be reserved for each CoS during congestion. This setting allows the physical Ethernet adapter to be shared by multiple vNICs and helps ensure a minimum bandwidth during congestion.

You should create separate classes of service for storage, live migration, cluster, host management and virtual machine data. Table 1 shows an example in which data from the different vNICs is segmented into different classes.

Table 1 Sample QoS Classification and Bandwidth Reservation.

Table 1. Virtual Machine Data High Priority

Traffic Type	COS	Minimum Bandwidth (Gbps)
Host Management and Cluster	2	10
Live Migration	1	10
Storage	3	40
Virtual Machine Data Low Priority	4	10
Virtual Machine Data High Priority	5	30

For all virtual machine interfaces (vEths) connected to the Cisco Nexus 1000V, the classification and CoS marking must be configured on the Cisco Nexus 1000V. The Cisco Nexus 1000V offers detailed QoS classification (based on Layer 2, Layer 3, and Layer 4 headers), allowing the administrator to prioritize certain application flows. This

level of visibility is not offered by the Cisco UCS VIC. For more information about configuring QoS classification and marking, refer to the Cisco Nexus 1000V QoS configuration guide.

Upstream Switch Connectivity

The Cisco Nexus 1000V can be connected to any upstream switch (any Cisco switch as well as switches from other vendors) that supports standards-based Ethernet and does not require any additional capability to be present on the upstream switch to function properly. Much of the design work for a Cisco Nexus 1000V solution focuses on proper upstream switch connectivity.

You can connect a Cisco Nexus 1000V Switch to a physical infrastructure using either of two means:

- **Standard uplinks:** A standard uplink is an uplink that is not a member of a PortChannel from the VEM to a physical switch. It provides no capability to load balance across multiple standard uplink links and no high-availability characteristics. When a standard uplink fails, no secondary link exists to take over. Defining two standard uplinks to carry the same VLAN is an unsupported configuration. Cisco NX-OS will generate a syslog message when this condition is detected.

Given the requirements of most data center networks, standard uplinks should rarely, if ever, be used in a Cisco Nexus 1000V design.

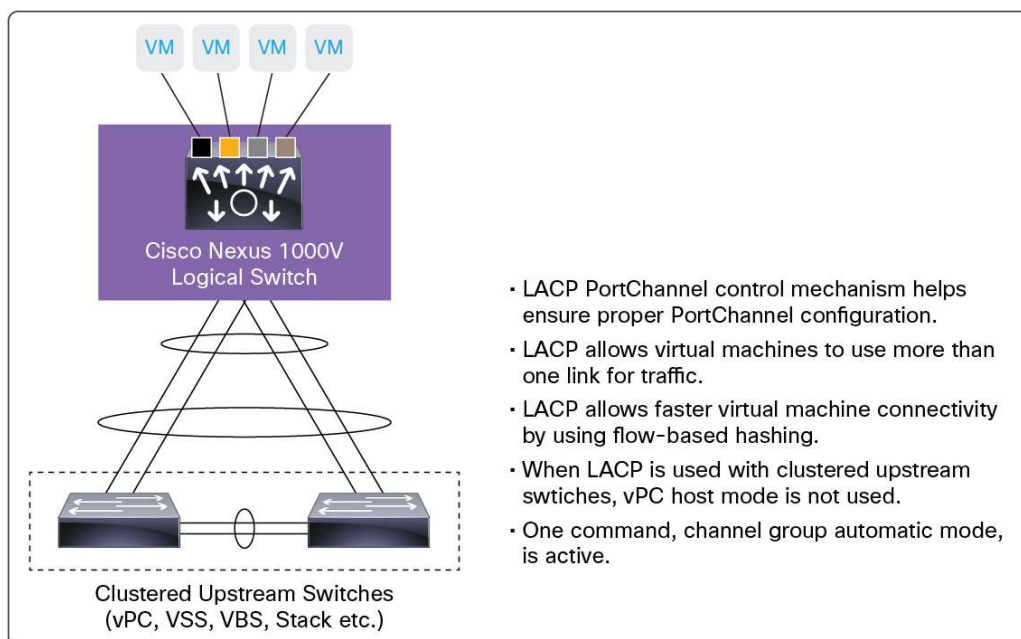
- **PortChannels:** Because Cisco Nexus 1000V Switches are end-host switches, the network administrator can use a different approach than can be used on physical switches, implementing a PortChannel mechanism in either of two modes:
 - **Standard PortChannel:** The PortChannel is configured on both the Cisco Nexus 1000V Switch and the upstream switch on the same number of ports. This approach is exactly the same as for regular PortChannel (EtherChannel) configuration on physical switches.
 - **Special PortChannel:** For some special PortChannels, such as virtual PortChannel host mode (vPC-HM) subgroups using Cisco Discovery Protocol or manual mode, no PortChannel configuration is also required on the upstream switch.

Regardless of the mode, PortChannels are managed using the standard PortChannel CLI construct, but each mode behaves differently.

Standard PortChannel

A standard PortChannel on Cisco Nexus 1000V Switches behaves like an EtherChannel on other Cisco switches and supports Link Aggregation Control Protocol (LACP). Standard PortChannels require that all uplinks in the PortChannel be in the same EtherChannel on the upstream switch (Figure 34).

Figure 34. Standard PortChannel Configuration



Standard PortChannels can be spread across more than one physical switch if the physical switches are clustered. Examples of clustered switching technology include the Cisco Catalyst® 6500 Virtual Switching System 1440, virtual PortChannels on the Cisco Nexus 7000 Series Switches, and the Cisco Catalyst Blade Switch 3120 for HP. Clustered switches act as a single switch and therefore allow the creation of EtherChannels across them. This clustering is transparent to the Cisco Nexus 1000V. When the upstream switches are clustered, the Cisco Nexus 1000V Switch should be configured to use LACP with one port profile, using all the available links. This configuration will make more bandwidth available for the virtual machines and accelerate Live Migration.

LACP Offload

Traditionally, LACP is processed on the control plane of a switch (the supervisor or VSM). Because the Cisco Nexus 1000V has a distributed model in the virtualization environment, new problems may arise when LACP is run on the VSM in this way.

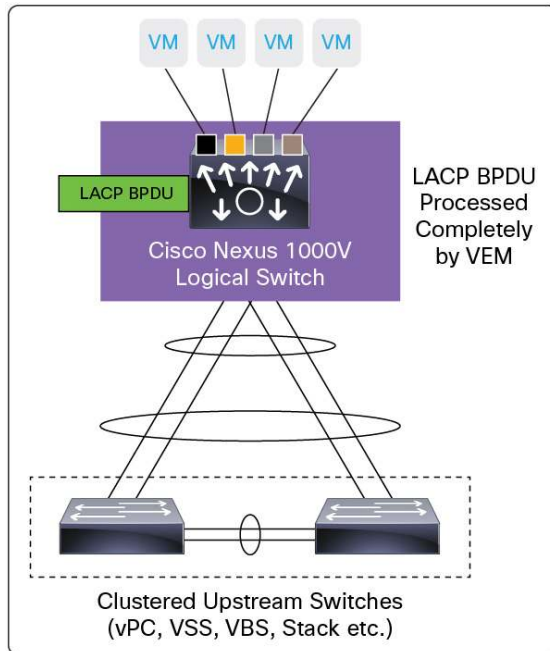
For example, if the VEM operates without the presence of the VSM, also referred to as headless mode, whenever a link flap or reboot of the VEM occurs and there is no communication between the VEM and VSM, an LACP bundle will not form. The bundle will not form because the VSM is responsible for originating and processing the LACP control packets, also referred to as LACP protocol data units (PDUs), needed to negotiate an LACP PortChannel.

Another situation with a similar problem may occur if the VSM is hosted (as a virtual machine) behind the VEM and remote storage is used in an FCoE deployment. In an FCoE deployment, the virtual Fibre Channel (vFC) interface is bound to the Ethernet LACP PortChannel.

For the VSM virtual machine to boot, this vFC interface needs to be up so that the remote storage can be accessed. However, because the VSM is responsible for originating and processing the LACP PDUs, a PortChannel cannot be formed because the VSM is not up at the start.

Beginning with Cisco NX-OS Release 5.2(1)SV1(5.1) for the Cisco Nexus 1000V Switch for Microsoft Hyper-V, LACP processing can be moved from the control plane (VSM) to the data plane (VEM), as shown in Figure 35.

Figure 35. LACP Offloaded to VEM in an FCoE Deployment



The origination and processing of LACP PDUs is now completely offloaded to the VEM. Therefore, you should enable this feature in any of the scenarios described in the preceding paragraphs, so that the VEM has no dependence on VSM connectivity while negotiating an LACP PortChannel. This feature makes the overall deployment of the Cisco Nexus 1000V solution more robust.

Special PortChannel

Most access-layer switches do not support clustering technology, yet most Cisco Nexus 1000V designs require PortChannels to span multiple switches. The Cisco Nexus 1000V offers several ways to connect the Cisco Nexus 1000V Switch to upstream switches that cannot be clustered. To enable this spanning of switches, the Cisco Nexus 1000V provides a PortChannel-like method that does not require configuration of a PortChannel upstream. There are two main vPC-HM configurations:

- vPC-HM MAC address pinning
- vPC-HM subgroups

vPC-HM MAC Address Pinning

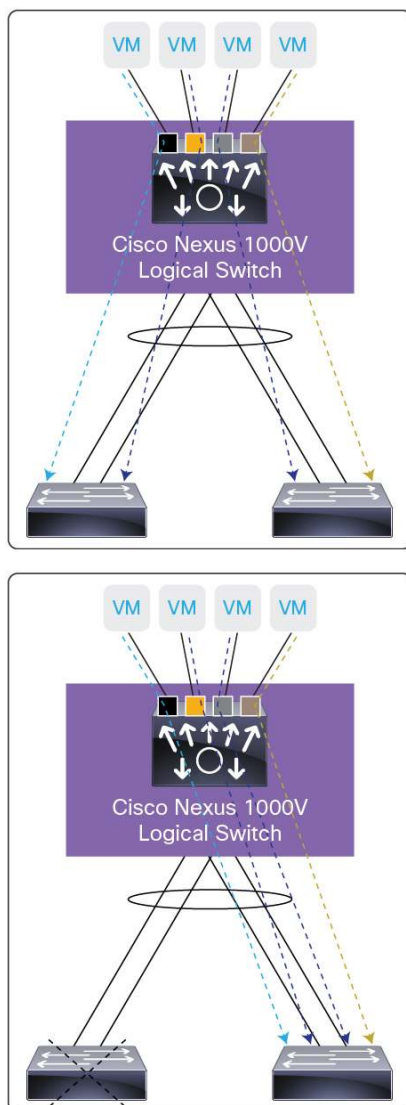
MAC address pinning defines all the uplinks coming from the server as standalone links and pins different MAC addresses to those links in a round-robin fashion. This approach helps ensure that the MAC address of a virtual machine is never seen on multiple interfaces on the upstream switches. Therefore, no upstream configuration is required to connect the Cisco Nexus 1000V VEM to the upstream switches (Figure 36).

Furthermore, MAC address pinning does not rely on any protocol to distinguish the various upstream switches, making the deployment independent of any hardware or design.

However, this approach does not prevent the Cisco Nexus 1000V Switch from constructing a PortChannel on its side, providing the required redundancy in the data center in the event of a failure. If a failure occurs, the Cisco Nexus 1000V Switch will send a gratuitous ARP packet to alert the upstream switch that the MAC address of the VEM learned on the previous link will now be learned on a different link, enabling failover in less than a second.

MAC address pinning enables consistent and easy deployment of the Cisco Nexus 1000V Switch because it does not depend on any physical hardware or any upstream configuration, and it is the preferred method for deploying the Cisco Nexus 1000V if the upstream switches cannot be clustered.

Figure 36. MAC Address Pinning



vPC-HM Subgroups

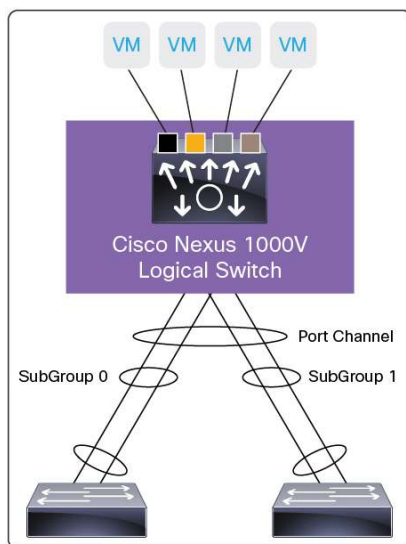
vPC-HM subgroups provide another way of creating a PortChannel on the Cisco Nexus 1000V Switch side when upstream switches cannot be clustered. With vPC-HM, the PortChannel configured on the Cisco Nexus 1000V Switch is divided into subgroups, or logical, smaller PortChannels, with each subgroup representing one or more uplinks to one upstream physical switch (Figure 37).

Links within the PortChannel that are connected to the same physical switch are bundled in the same subgroup automatically through the use of the Cisco Discovery Protocol packets received from the upstream switch. Alternatively, interfaces can be manually assigned to a specific subgroup using interface-level configuration.

When vPC-HM is used, each vEth interface on the VEM is mapped to one of the two subgroups using a round-robin mechanism. All traffic from the vEth interface uses the assigned subgroup unless the assigned subgroup is unavailable, in which case the vEth interface will fail over to the remaining subgroup. When the originally assigned subgroup becomes available again, traffic will shift back to its original location. Traffic from each vEth interface is then hashed within its assigned subgroup based on the configured hashing algorithm.

When multiple uplinks are attached to the same subgroup, the upstream switch needs to be configured in a PortChannel, bundling those links together. The PortChannel needs to be configured with the option mode on.

Figure 37. Virtual PortChannel Host Mode



Load Balancing

The Cisco Nexus 1000V provides 17 hashing algorithms to load-balance traffic across physical interfaces in a PortChannel. These algorithms can be divided into two categories: source-based hashing and flow-based hashing. The type of load balancing that the Cisco Nexus 1000V uses can be specified in VEM-level detail, so one VEM can implement flow-based hashing, using the better load sharing offered by that mode, and another VEM not connected to a clustered upstream switch can use MAC address pinning and thus source-based hashing.

Choose the hashing algorithm used with care because it affects the available configuration options and may require configuration changes on the access-layer switches. The default hashing algorithm used by the Cisco Nexus 1000V is source MAC address hashing (a source-based hash).

Source-Based Hashing

Source-based hashing algorithms help ensure that a MAC address is transmitted down only a single link in the PortChannel, regardless of the number of links in a PortChannel.

With source-based hashing, a MAC address can move between interfaces under the following conditions:

- The virtual machine moves to a new Microsoft Hyper-V A link fails, causing recalculation of the hashing.

The following Cisco Nexus 1000V algorithms can be classified as source-based hashes:

- Virtual port ID
- Source MAC address

Flow-Based Hashing

Flow-based hashing enables traffic from a single MAC address to be distributed down multiple links in a PortChannel simultaneously. Use of a flow-based hash increases the bandwidth available to a virtual machine or to Live Migration and increases the utilization of the uplinks in a PortChannel by providing more precise load balancing.

Flow-based hashing algorithms are any algorithms that use the following to hash:

- Packet destination
- Layer 4 port
- Combinations of source address, destination address, and Layer 4 port

Network-State Tracking

Network-state tracking (NST) is a mechanism that is used to detect Layer 1 and network connectivity failures that would otherwise cause virtual machine traffic to be dropped during an uplink failure. An uplink failure occurs when the upstream switch encounters a driver or firmware failure that prevents it from signaling a link to the neighboring link interface (the Cisco Nexus 1000V VEM in this case). NST mitigates this problem by using NST packets, which probe interfaces on other subgroups of the same VEM.

When a link failure occurs on the upstream switch, NST can detect the failure by sending the tracking packet from one interface in Subgroup 0 to all interfaces in Subgroup 1. Because a link failure occurred on the upstream switch, this tracking packet will not be received by Subgroup 1. The vPC-HM PortChannel interface will then be identified as split, and a syslog will be generated for this split. Also, packet counters are monitored to detect whether there is traffic coming into the subgroup that did not receive the tracking packet. If packet counters are not incrementing, then the upstream switch is not transmitting to the VEM in the subgroup. At this point, traffic that was originally pinned to the Subgroup 1 interface will be repinned to the active interface in Subgroup 0.

You should enable NST when an uplink failure may cause traffic to be dropped. The most common deployment of NST is in cases in which third-party networking entities may encounter driver or firmware or other software failures in which the interface is not brought down completely and the VEM interface continues to forward traffic.

Cisco Nexus 1000V Switch for Microsoft Hyper-V Sample Configuration

This section presents a sample Cisco Nexus 1000V Switch for Microsoft Hyper-V configuration.

1. Create logical network and network segment pools.

This example creates a logical network called DMZ with two network segment pools - DMZ-NY and DMZ-SFO - to model a DMZ network.

```
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# nsm logical network DMZ
Nexus1000V(config-logical-net)# description Logical Network for DMZ VMs
Nexus1000V(config-logical-net)# exit
Nexus1000V(config)#
Nexus1000V(config)#

Nexus1000V(config)#
Nexus1000V(config)# nsm network segment pool DMZ-SFO
Nexus1000V(config-net-seg-pool)# member-of logical network DMZ
Nexus1000V(config-net-seg-pool)# exit
Nexus1000V(config)# nsm network segment pool DMZ-NY
Nexus1000V(config-net-seg-pool)# member-of logical network DMZ
Nexus1000V(config-net-seg-pool)# exit
Nexus1000V(config)#
```

2. Create an IP pool template.

Microsoft SCVMM can behave like a DHCP server, allocating IP addresses to virtual machines. When virtual machines are deployed from the Microsoft SCVMM console, Microsoft SCVMM can set the IP address and the default gateway on the virtual machines. When Cisco Nexus 1000V is used to manage virtual networking on Microsoft Hyper-V, the network administrator must define the IP-pool range to be used when virtual machines are deployed on a VLAN-based virtual machine network. This example creates an IP pool for virtual machines deployed on Vlan10.

```
Nexus1000V(config)#
Nexus1000V(config)# nsm ip pool template Vlan10IPPool
Nexus1000V(config-ip-pool-template)# ip address 20.20.20.50 20.20.20.70
Nexus1000V(config-ip-pool-template)# network 20.20.20.1 255.255.255.0
Nexus1000V(config-ip-pool-template)# default-router 20.20.20.1
Nexus1000V(config-ip-pool-template)# exit
Nexus1000V(config)#
```

3. Create a network segment.

This example creates a VLAN-backed network segment called Vlan10. This network segment is placed in the DMZ-SFO network segment pool.

```
Nexus1000V(config)#
Nexus1000V(config)# nsm network segment Vlan10
Nexus1000V(config-net-seg)# member-of network segment pool DMZ-SFO
Nexus1000V(config-net-seg)# switchport access vlan 10
Nexus1000V(config-net-seg)# switchport mode access
Nexus1000V(config-net-seg)# ip pool import template Vlan10IPPool
Nexus1000V(config-net-seg)# publish network segment
Nexus1000V(config-net-seg)# exit
Nexus1000V(config)#
```


4. Create a virtual Ethernet port profile with an IP access list.

This example creates a simple IP access list that denies IP traffic from virtual machine 20.20.20.22 to 20.20.20.20 and applies this access list to a vEth port profile.

```
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# ip access-list Restrict-ContractorVM
Nexus1000V(config-acl)# 10 deny ip 20.20.20.22/32 20.20.20.20/32
Nexus1000V(config-acl)# 20 permit ip any any
Nexus1000V(config-acl)# exit
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# port-profile type vethernet RestrictedProfile
Nexus1000V(config-port-prof)# ip port access-group Restrict-ContractorVM in
Nexus1000V(config-port-prof)# no shutdown
Nexus1000V(config-port-prof)# state enabled
Nexus1000V(config-port-prof)# publish port-profile
Nexus1000V(config-port-prof)# exit
Nexus1000V(config)#
Nexus1000V(config)#
```

5. Create an Ethernet port profile.

This example creates an Ethernet port profile that uses vPC host mode for redundancy. Typically, all Ethernet port profiles are configured as trunk ports.

```
Nexus1000V(config)#
Nexus1000V(config)# port-profile type ethernet UplinkProfile
Nexus1000V(config-port-prof)# channel-group auto mode on mac-pinning
Nexus1000V(config-port-prof)# no sh
Nexus1000V(config-port-prof)# state enabled
Nexus1000V(config-port-prof)# switchport mode trunk
Nexus1000V(config-port-prof)# switchport trunk allowed vlan 10
Nexus1000V(config-port-prof)# switchport trunk native vlan 10
Nexus1000V(config-port-prof)# exit
Nexus1000V(config)#
Nexus1000V(config)#
```

6. Create an uplink network.

The **network uplink** command is new in the Cisco Nexus 1000V Switch for Microsoft Hyper-V. Each uplink network configured on the VSM is available as an uplink port profile to the Microsoft SCVMM administrator. The example here creates an uplink network that uses the Ethernet profile UplinkProfile and allows the network segment pools: DMZ-SFO and DMZ-NY. The uplink network defines the network policy applied to switch uplink interfaces and the VLANs allowed on them.

```
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# nsm network uplink Nexus1000VUplinkProfile
Nexus1000V(config-uplink-net)# import port-profile UplinkProfile
Nexus1000V(config-uplink-net)# allow network segment pool DMZ-NY
Nexus1000V(config-uplink-net)# allow network segment pool DMZ-SFO
Nexus1000V(config-uplink-net)# publish network uplink
Nexus1000V(config-uplink-net)# exit
Nexus1000V(config)#
Nexus1000V(config)#
Nexus1000V(config)# copy r s
[#####] 100%
Copy complete, now saving to disk (please wait)...
Nexus1000V(config)#
```

Conclusion

Cisco Nexus 1000V Switches integrate into the expanding virtualized data center, providing secure, nondisruptive, policy-based enhanced networking features and visibility for the networking team into the server virtualization environment. The comprehensive feature set of the Cisco Nexus 1000V allows the networking team to troubleshoot more rapidly any problems in the server virtualization environment, increasing the uptime of virtual machines and protecting the applications that propel the data center.

For More Information

For more information about the Cisco Nexus 1000V Switch, please refer to the following:

- Cisco Nexus 1000V Switch product information: <http://www.cisco.com/en/US/products/ps13056/index.html>.
- Cisco Nexus 1000V Switch installation and upgrade documentation: http://www.cisco.com/en/US/products/ps13056/tsd_products_support_install_and_upgrade.html.
- Cisco Nexus 1000V Switch installation and configuration documentation: http://www.cisco.com/en/US/products/ps13056/products_installation_and_configuration_guides_list.html.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)