Cisco Nexus 1000V on HP BladeSystem

An informational guide for optimal VMware vSphere design using HP BladeSystem (w/Virtual Connect) and Cisco Nexus network infrastructure

Table of contents

Overview

The introduction of virtual machine technology changes the way servers and networks are designed. A physical server can now contain multiple "virtual" servers communicating with multiple "virtual" NICs. In addition, the network infrastructure is now extended beyond the physical NIC into the confines of the physical server itself. This is accomplished by the introduction "virtual" switches also contained within the physical server. All of these new virtual components essentially create virtual network infrastructure; all self-contained within a physical server.

The physical connectivity to servers capable of running virtual machines continues to evolve. In 2007, HP introduced Virtual Connect to the BladeSystem family. Virtual Connect defines a physical "end-point" (at the Virtual Connect uplink port) to the customer network infrastructure and allows a server administrator to control the assignment of networks (or VLANS) to physical NICS inside the BladeSystem enclosure. This VC "end point" cleanly separates the role of the server administrator from the role of the network administrator. In addition, Virtual Connect allows the server administrator to manage the assignment of networks and physical MAC addresses to physical NICs on the blades contained in the BladeSystem enclosure. All of these assignments are made by assigning a Virtual Connect "Server Profile" to a BladeSystem slot. Whatever blade is inserted into the slot, gains the characteristics of the server profile assigned to that slot.

Server virtualization on BladeSystem hardware (leveraging Virtual Connect) has become very popular and VMware has emerged as a leader and innovator in the server virtualization marketplace. Furthermore, VMware has continued this evolution with the release of VMware vSphere 4 in April 2009. One of the new features of VMware vSphere 4 is the introduction of the new vNetwork Distributed Switch (vDS). The vDS extends the feature set of the VMware Standard Switch through an abstracted, single distributed switch representation of multiple ESX servers.

In another recent innovation, VMware began shipping a virtual switch API. This API enables 3rd-party vendors to distribute their own virtual switch. Cisco has developed the Cisco Nexus 1000V virtual switch. The 1000V allows a Cisco network administrator to fully manage the virtual switch inside the physical server using Cisco tools. While the 1000V enables many new virtual switch capabilities, the ramifications of enabling some of these options can have serious consequences for the rest of the network infrastructure.

The purpose of this document is to review the pros/cons of using the various capabilities of the Cisco Nexus 1000V virtual switch and discuss the impact of the 1000V to the surrounding network and Virtual Connect infrastructures.

VMware Virtual Switching Background

VMware vSphere-based vNetwork offers three different virtual switch options for virtual networking:

1. Standard Virtual Switch
2. Distributed Virtual Switch
3. Third Party Distributed Virtual Switch

This section briefly describes each of these virtual switch technologies. If more information is required, please refer to the Virtual Networking Concepts document at: http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf.

1. Standard Virtual Switch Architecture[1]

The new vSphere-based standard Virtual Switch is functionally similar to the virtual switch included with VI3 (ESX 3.x). VI3 network configuration is done at the host level. Each virtual machine and the service console have one or more virtual network adapters, or vNICs. The operating system and applications communicate to a vNIC through a standard device driver or a VMware optimized device driver. The VMkernel also has vNICs for VMotion and IP storage network requirements.

Each ESX host has its own virtual switches (AKA "Standard Switch"). On the downlink side of the virtual switch are port groups that connect to virtual machines. On the other side of the virtual switch are uplink connections to physical Ethernet adapters on the server where the virtual switch resides. Virtual machines, the service console, and VMkernel components connect to the outside world through the physical Ethernet adapters that are connected to the virtual switch uplinks.

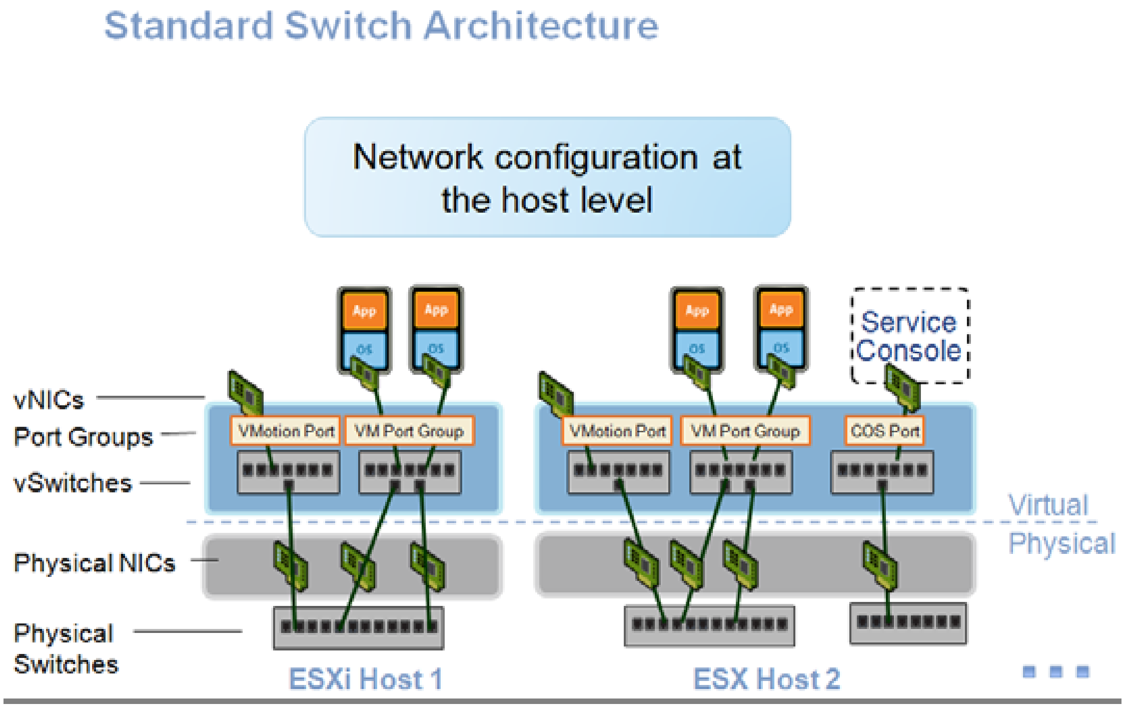

[1] *Excerpts taken from the VMware publication: VMware Virtual Networking Concepts.*

2. Distributed Virtual Switch Architecture[2]

A distributed virtual switch is a managed entity that is configured inside vCenter. Distributed virtual switches provide the same basic functions as standard vSwitches, but they exist across two or more clustered ESX or ESXi hosts. vCenter Server owns the configuration of distributed switches, and the configuration is consistent across all hosts.

Like a VI3 standard vSwitch, a distributed switch connects to a physical network via one or more physical Ethernet adapters on the hosts included in the cluster. In this manner, physical NICs become clustered resources to use as required by the networked components. Each distributed switch includes distributed ports. A distributed port represents a port to which you can connect any networking entity, such as a virtual machine, the Service Console, etc. vCenter Server stores the state of distributed ports in the vCenter database; allowing networking statistics and policies migrate with virtual machines when moved from host to host.

A distributed switch is not the same as a single switch spanning across several hosts. Two virtual machines on different hosts can communicate with each other only if both virtual machines have uplinks in the same broadcast domain (or have a route established). Think of a distributed switch as a template for the network configuration on each ESX or ESXi host.

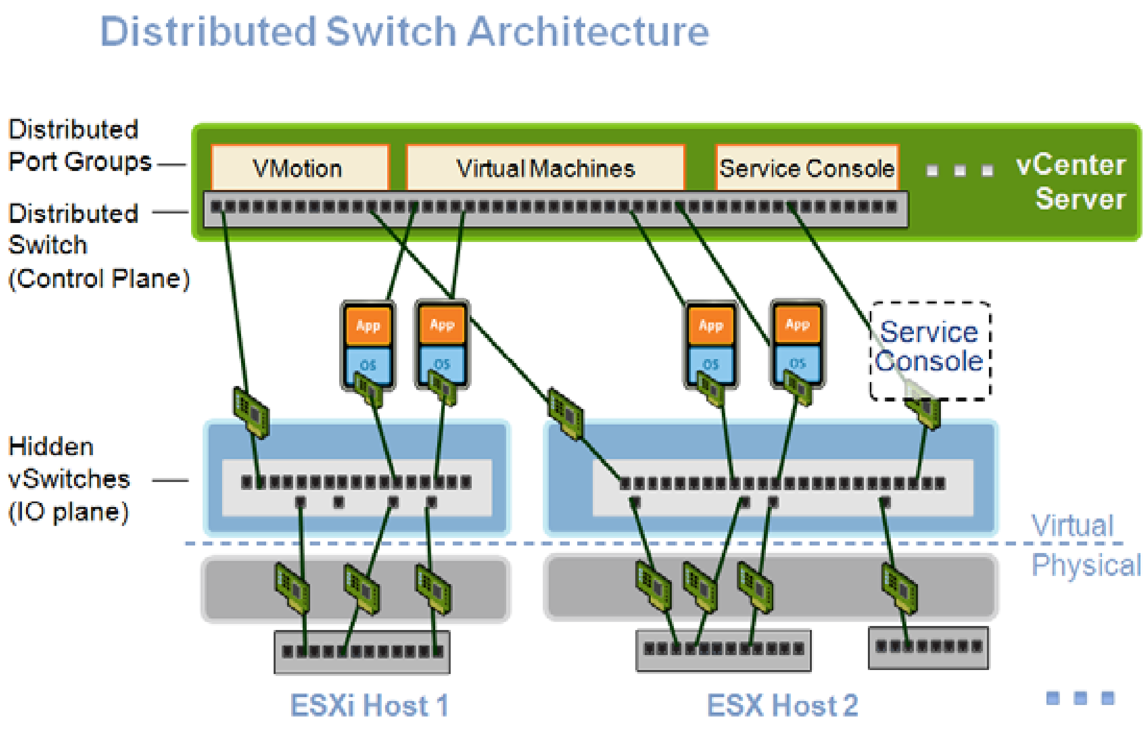

[2] *Excerpts taken from the VMware publication: VMware Virtual Networking Concepts.*

3. Third Party Distributed Virtual Switches[3]

VMware vSphere introduced vNetwork Appliance APIs. These APIs allow third-party developers to create distributed switch solutions for use in a VMware Virtual Infrastructure datacenter.

The diagram below shows the way a third-party solution plugs in to the vNetwork architecture. The components are:

- A Custom Control Plane is implemented outside of vCenter, for example it may be implemented as a virtual appliance.
- VI Client includes a plug-in to provide a management interface.
- vCenter includes an extension to handle the communication with the control plane.
- On the host, a custom IO plane agent replaces the standard IO plane agent.
- And the IO plane itself may be replaced for customization of forwarding and filtering.

The Cisco Nexus 1000V is the first third-party switch to leverage vNetwork Appliance APIs. Network administrators can use this solution in place of the vNetwork Distributed Switch. This document explores the 1000V in great detail.

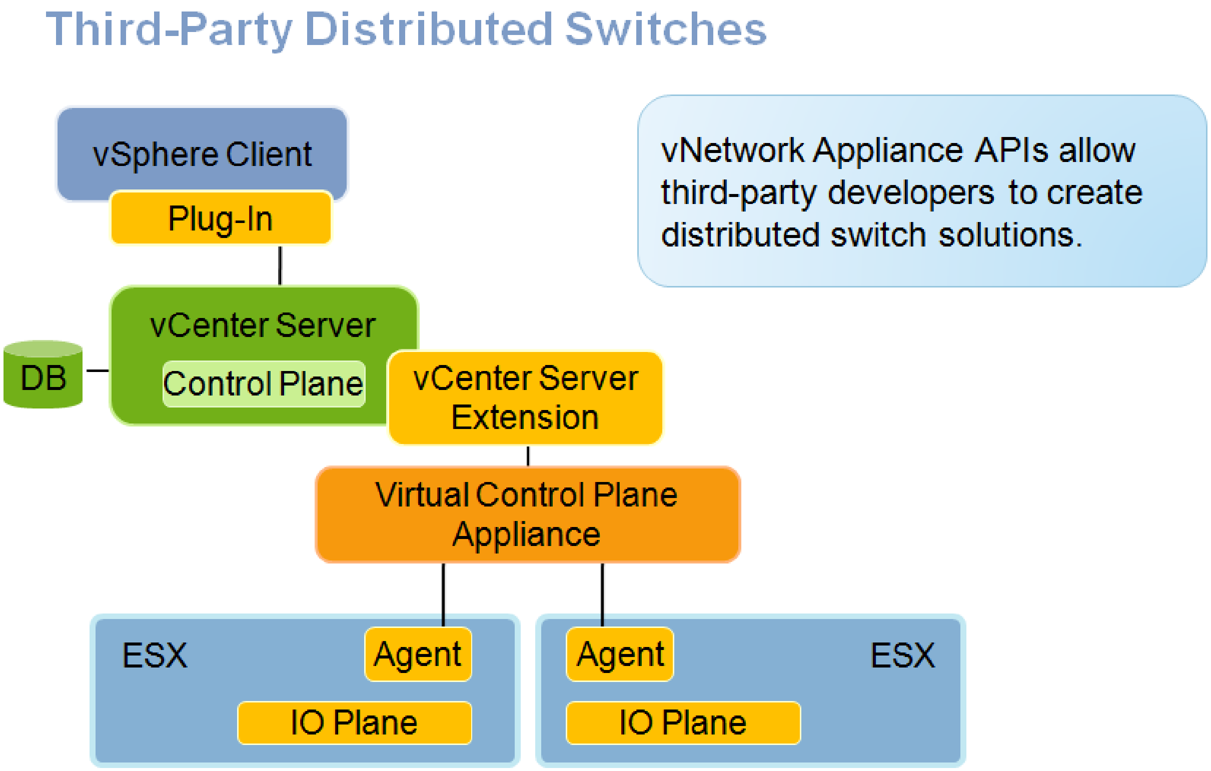

Third-Party Distributed Switches

vNetwork Appliance APIs allow third-party developers to create distributed switch solutions.

[3] *Excerpts taken from the VMware publication: VMware Virtual Networking Concepts.*

License Requirements for the vDS and Cisco Nexus 1000V

With the introduction of vSphere 4.0, VMware has introduced a new tiered license model. Customers that standardized on the VI3 Enterprise license model will be unable to take advantage of the full vSphere product offering without enduring additional upgrade costs. Existing VI3 Enterprise customers with current Support-and-Subscription (SnS) will be able to upgrade and migrate to vSphere 4.0 Enterprise. However, vSphere 4.0 Enterprise does not include one of VMware's newest technologies: the vNetwork Distributed Switch (vDS). In order to deploy and utilize VMware's vDS, a vSphere Enterprise customer will be required to upgrade to vSphere Enterprise Plus.

The Cisco Nexus 1000V is the first third party distributed virtual switch offered for vSphere. As a distributed virtual switch, the 1000V requires vSphere Enterprise Plus licensing as well as a separate license for the 1000V itself. The 1000V license is based on per socket per host model and incurs additional support costs.

Table # outlines the entire vSphere license tiers and their licensed features.

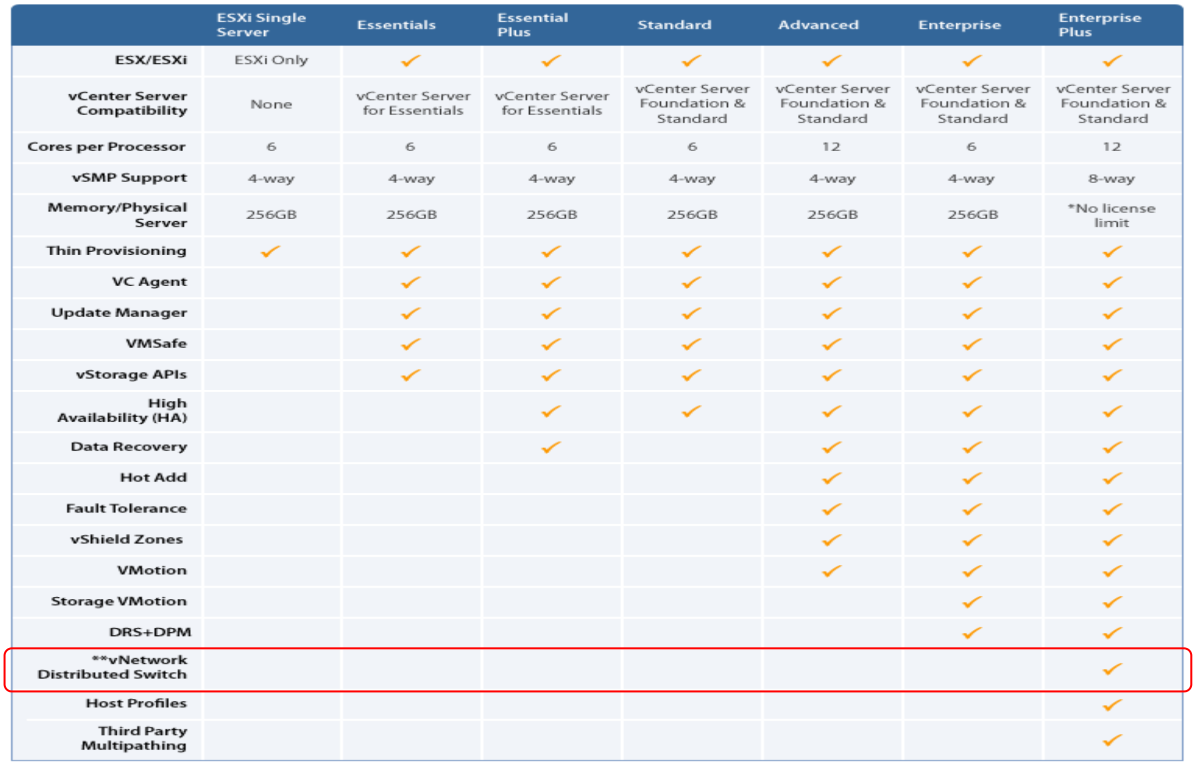

	ESXi Single Server	Essentials	Essential Plus	Standard	Advanced	Enterprise	Enterprise Plus
ESX/ESXi	ESXi Only	✓	✓	✓	✓	✓	✓
vCenter Server Compatibility	None	vCenter Server for Essentials	vCenter Server for Essentials	vCenter Server Foundation & Standard	vCenter Server Foundation & Standard	vCenter Server Foundation & Standard	vCenter Server Foundation & Standard
Cores per Processor	6	6	6	6	12	6	12
vSMP Support	4-way	4-way	4-way	4-way	4-way	4-way	8-way
Memory/Physical Server	256GB	256GB	256GB	256GB	256GB	256GB	*No license limit
Thin Provisioning	✓	✓	✓	✓	✓	✓	✓
VC Agent		✓	✓	✓	✓	✓	✓
Update Manager		✓	✓	✓	✓	✓	✓
VMSafe		✓	✓	✓	✓	✓	✓
vStorage APIs		✓	✓	✓	✓	✓	✓
High Availability (HA)			✓	✓	✓	✓	✓
Data Recovery			✓		✓	✓	✓
Hot Add					✓	✓	✓
Fault Tolerance					✓	✓	✓
vShield Zones					✓	✓	✓
VMotion					✓	✓	✓
Storage VMotion						✓	✓
DRS+DPM						✓	✓
**vNetwork Distributed Switch							✓
Host Profiles							✓
Third Party Multipathing							✓

** VMware ESX4.0 and ESXi 4.0 currently provide technical support for up to 1TB of memory. Review the vSphere 4 Configuration Maximums document for more information.*

*** Enabler for 3rd party switch support. May require an additional license cost from 3rd party vendor.*

Technical Summary of New Technologies

Cisco Nexus 1000V

The Cisco Nexus 1000V is the first 3rd party distributed virtual switch (vDS) that can be used by vSphere ESX to enable Cisco-based networking functions within the ESX host. The 1000V is based on the Cisco NX-OS data center operating system. Common networking functions for monitoring, administration, security, and provisioning are provided with this virtual switch.

The 1000V has two main components, the Virtual Ethernet Module (VEM) and the Virtual Supervisor Module (VSM).

Virtual Ethernet Module (VEM)

The VEM manifests itself as software installed in the vSphere ESX hypervisor. It enables networking and security features, performs switching between directly attached virtual machines, and provides uplink capabilities to the rest of the network infrastructure. Each VEM will forward packets independent of each other, there is no address learning across VEMs, no switching between VEMs, and no etherchannel across VEMs. The VEM replaces the embedded VMware vSwitch.

Virtual Supervisor Module (VSM)

The VSM is a control point responsible for the configuration, management, and monitoring of the 1000V systems. The VSM uses a version of the NX-OS and it can be managed through the NX-OS CLI. The VSM is integrated into vCenter for management as a standalone (physical) or virtual appliance. A single VSM can control and manage up to 64 VEM's as a single network device. The VSM can also be deployed as an active passive redundant pair.

The VSM can be deployed as a Virtual Appliance, a VM running on an ESX host, or installable via ISO or OVA file.

Current VSM deployment requirements:

- 2GB dedicated RAM (not shared), 1Ghz vCPU
- VSM should not be VMotioned

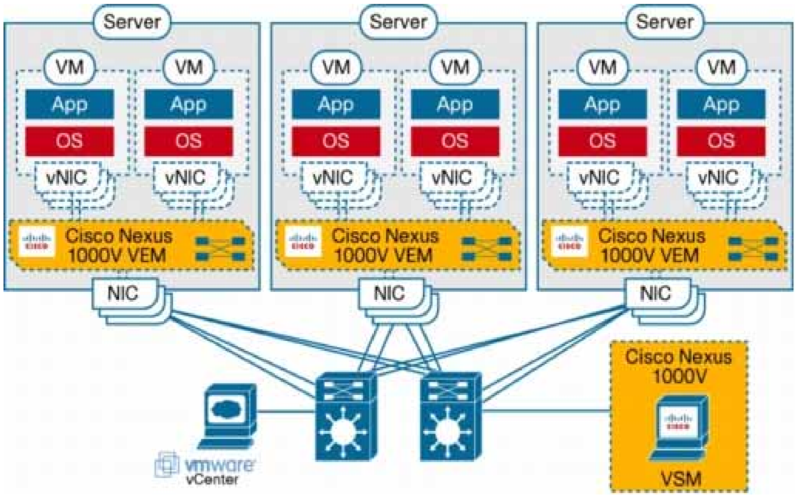

This illustration depicts the replacement of the native VMware distributed virtual switch (vDS) with the Nexus 1000V across 3 physical ESX hosts. It shows 3 VEMs and the required VSM management appliance.

Port Profiles

Another component of the Cisco Nexus 1000V is a construct known as a port profile. A Port Profile is an object or container that is used to define a set of common configuration commands that can be applied dynamically to either virtual or physical interfaces. A Port Profile can be defined once and deployed many times. Any changes made to an enabled Port Profile will cause a change in configuration to all interfaces using that profile. Port Profiles include definitions for port management, VLAN, PVLAN, Port-Channel, ACL, Netflow, Port Security, remote port mirroring, and QoS.

The Port Profile is managed and configured by the network administrator at the VSM. Once created, the Port Profile is pushed to the VMware vCenter. The port profile is displayed as a port group (within vCenter) that the virtual machine administrator can select when creating vNICs. When the Virtual Machine is powered on/off, the associated Port Profile used to dynamically configure the vEth in the VN-Link Switch. In other words, when the vm is turned on the port profile is responsible for creating the vEth connection to the 1000V.

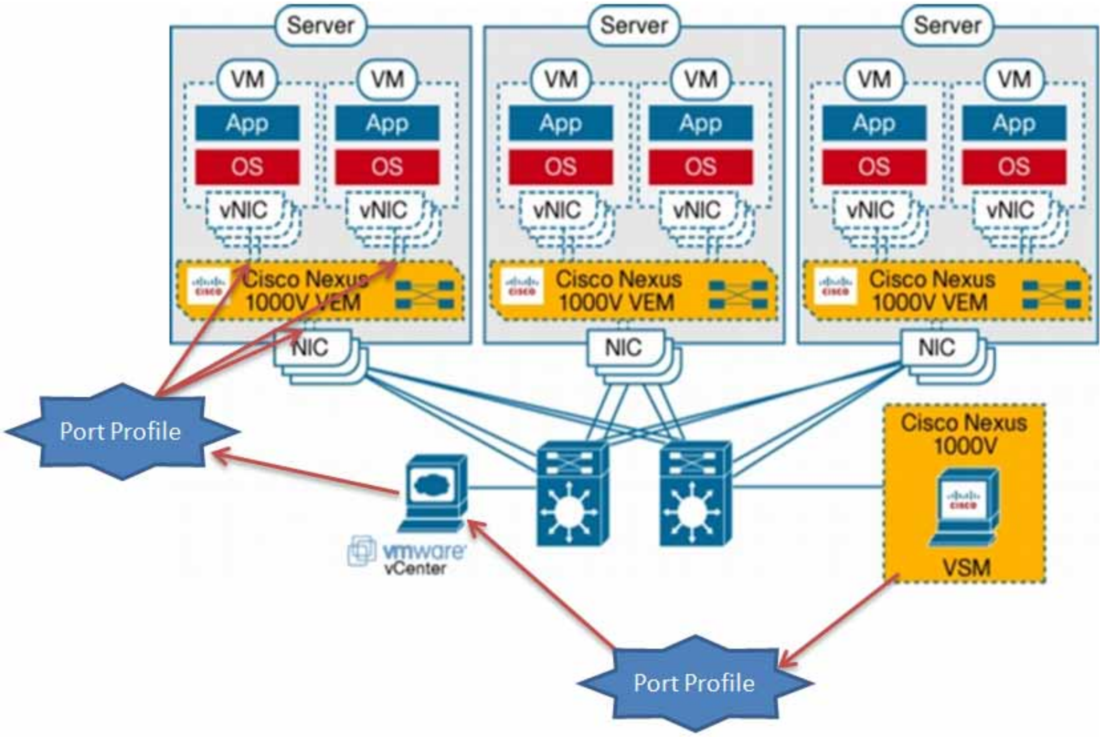

This illustration depicts creation of a Port Profile at the VSM, replication to the vCenter management application, followed by application to the vNICs in use by a given VM.

What is VN-Link?[4]

Using Cisco Terminology, the term VN-Link indicates the creation of a logical link between the vNIC on a virtual machine and a Cisco switch enabled for VN-Link. The mapping is the logical equivalent of using a cable to connect a NIC with a network port of a switch. The goal is to have features and capabilities that enable the virtual machine networking components (vNIC and vEth) to be managed like physical switch ports. These components are manifested as sort of a “virtual line card” in a VN-Link enabled Cisco switch.

Benefits: Policy Based VM Connectivity; Mobility of Network and security properties – Port Profiles; Ensures visibility and continued connectivity during VMotion; VN-Link works with the 1000V (software): Nexus 5000 with VN-Link (hardware)

Key Features: Virtual Ethernet Interfaces (vEth), these interfaces are managed the same as physical network access ports. These interfaces are created and stored on a VN-Link enabled switch. The vEth contains all attributes for configuration, security, and statistics for a virtual interface. The VN-Link enabled switch creates a logical mapping between the vEth interface and the corresponding vNIC on the VM. These vEth interfaces are transportable and can move from one physical server to another or from one physical access layer switch to another, enabling the VM to retain its defined network settings between hosts.

Port Profiles are the interface configuration parameters that are applied at the physical or virtual interfaces. The VN-Link enabled switch can push port profile information to the VMware vCenter management console where they are seen as distinct port groups and can be applied to vNIC interfaces.

VN-Link is implemented as a Cisco Distributed Virtual Switch within the hypervisor, Nexus 1000V, or as a new hardware device that has network interface virtualization enabled (VNTag).

Note: VN-Tagging requires that ALL switches between the VN-Link enabled switch and the ESX-based 1000V must support VN-Tag

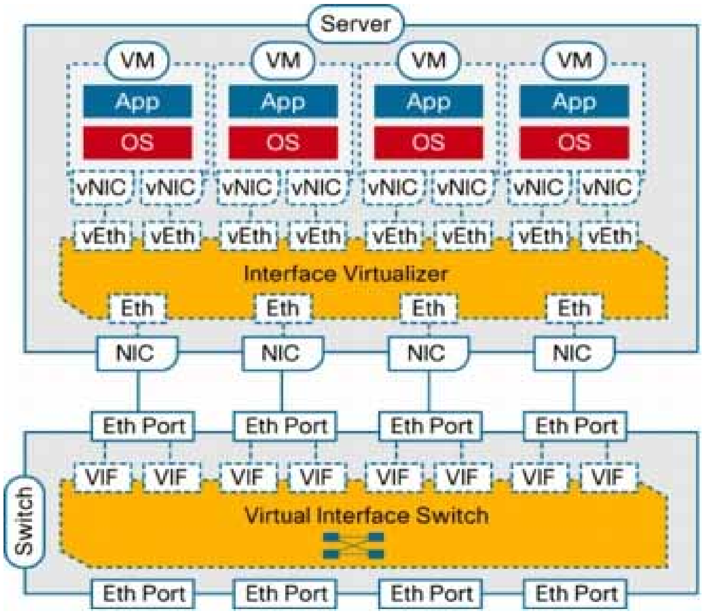

This illustration provides a logical depiction of a VN-Link enabled switch mapping to the vEth and vNIC interfaces of a 1000V VEM.

[4] *Some content taken from the Cisco publication: Cisco VN-Link: Virtualization-Aware Networking Technical Primer*

What is VNTag?

VNTag (Virtual Network Tag) is Cisco's proprietary implementation of a networking data frame header that can be used by Nexus switch devices. The VNTag communicates the identity of the port being presented (VLAN, QoS, Security, etc.). The VNTag enables a virtual machine vEth port to be logically attached to the network; allowing it to be fully managed like a physical Cisco port (or remote line card). This is commonly referred to as Network Interface Virtualization (NIV).

Use of a VNTag requires the decoupling of the switching functions from the hypervisor and relocates it in an external hardware-based network switch. As a result, the interface virtualizer or virtualization adapter never actually performs local switching between the virtual machines (inside the ESX host). Instead, switching is done by the NIV capable network switch (configured as the controlling bridge) that the interface virtualizer connects to. As a result, all frames entering a 1000V VEM supporting VNTag are egressed to an external controlling bridge.

VNTag is enabled by two design elements:

1. Hardware: Connecting a hardware based virtualization interface adapter (such as the Fabric Extender or the Cisco UCS adapter) inside the blade server to a Nexus based core switch (like the Nexus 5000 or UCS6100) that supports NIV. One of the core Nexus switches in this configuration must be configured as the controlling bridge.
2. Software: The Nexus OS (NX-OS) has the ability to enable VNTag by a configuration parameter.

As of the writing of this document, VNTag is not a ratified IEEE standard and VNTag is not available for use on the Cisco Nexus 1000V. Use of VN-Tag requires an external Cisco Fabric extender and Cisco Nexus switch. One of these Nexus switches must be configured as the controlling bridge. Currently, no details are provided on how to enable and provision VNTag on the Cisco Nexus 1000V.

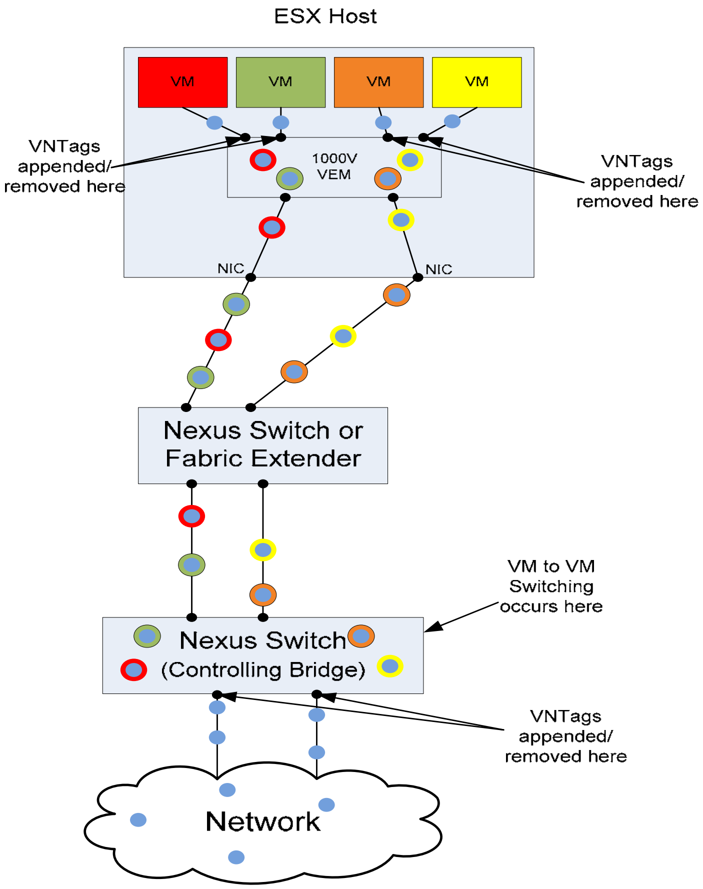
ESX Host
VM
VM
VM
VM
VNTags appended/ removed here
1000V VEM
VNTags appended/ removed here
NIC
NIC
Nexus Switch or Fabric Extender
VM to VM Switching occurs here
Nexus Switch (Controlling Bridge)
VNTags appended/ removed here
Network

This simplified illustration shows the 1000V VEM appending a VNTag to each frame of data and the necessary switching function being performed by the controlling bridge.

Furthermore, the use of VNTagging requires a frame to travel to the controlling bridge for all necessary switching functions. For instance, if a virtual machine wants to communicate from VM (red) to VM (orange) on the same ESX host (as shown below), the data frame will egress the 1000V (and ESX host) traveling to the controlling bridge and back down to the destination VM. This process can introduce significant bandwidth demands and limit the total bandwidth of the vSwitch to the bandwidth that the physical uplink is configured to, i.e.: a single 10Gb uplink

When VNTagging is not in use and a virtual machine needs to communicate from VM (red) to VM (orange) on the same ESX host (as shown below) the data frames are allowed to stay local to the ESX host; as the virtual switch VEM passes them directly to the virtual Ethernet port for the destination VM. Without VNTagging, communication from VM (red) to VM (orange) is extremely efficient (but local). Only traffic destined for systems outside of the vSwitch will travel out the uplink.

See the drawings on the next page for a visual comparison of the data frame flow with VNTagging vs. without VNTagging support.

1000V Communication Path (VNTag)

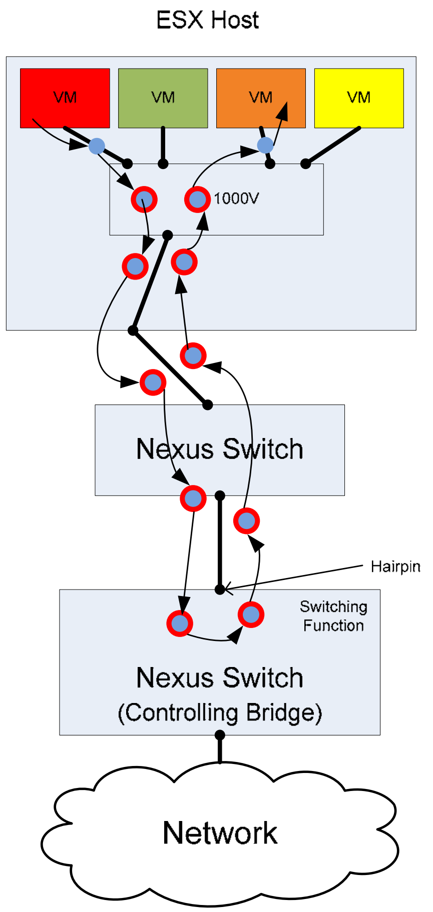

1000V Communication Path (no VNTag)

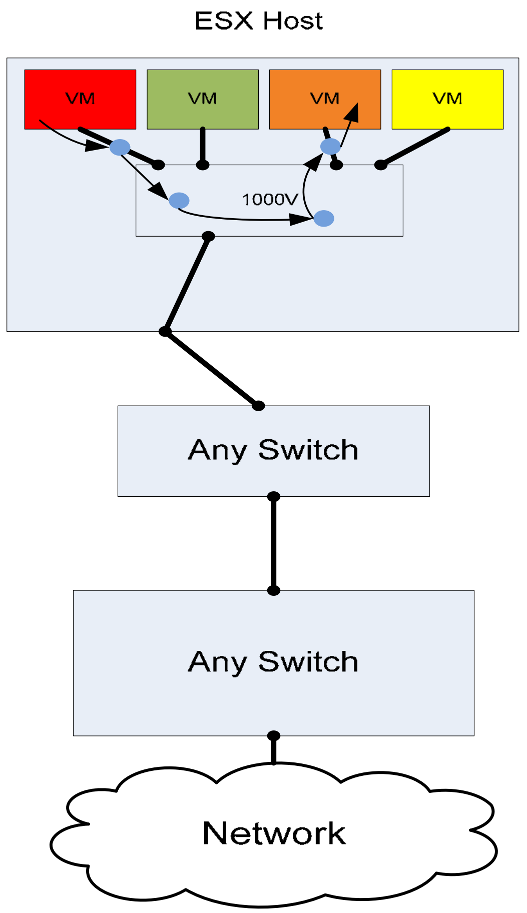

On the surface, the use of VNTag may sound inefficient, but there are pros and cons to its use. The use of VNTag enables a network administrator to manage virtual and physical ports as a remote line card in the controlling bridge. While this is nice from an operational management perspective, it requires the use of a proprietary "tag" that is only understood by Cisco Nexus technology. Secondly, all VNTagged frames must be switched at the controlling bridge. As a result, as frames move from the edge devices up towards the controlling bridge, additional bandwidth is required. If all frames move to the controlling bridge, the bandwidth requirements of the devices closer to the controlling bridge increase dramatically.

Choosing the 1000V - Impact to Customers

The Cisco Nexus 1000V can have a significant impact to a customer infrastructure. The 1000V offers customers a Cisco-branded virtual switch device deep within the ESX host that allows customers to fully manage their physical and virtual network switches using a consistent set of tools. In addition, the 1000V will offer customers a facility to force egress of virtual machine traffic into the network infrastructure. Both of these features are compelling to the network administrator. However, enablement of these features comes at a cost.

Acquisition Cost of the Cisco Nexus 1000V

The Cisco Nexus 1000V requires vSphere Enterprise Plus licensing to license the ESX host for distributed virtual switch technology. In addition, a separate license for the 1000V itself is required. The licensing model for the 1000V is based on the number of CPU sockets of the ESX servers attached as VEMs to the VSM.

Infrastructure Cost of Using VNTag on the Cisco Nexus 1000V

The use of VNTag capabilities requires a non-standard frame format that it incompatible with non-Nexus switches (including Cisco Catalyst switches). As a result, customers choosing to implement VNTag will be required to replace existing network switches with Cisco Nexus switches. All non-Nexus network devices located between the ESX host (hosting the 1000V) and the controlling bridge must be replaced.

Management & Support

Add-on components such as the Nexus 1000V that replace native distributed virtual switch functionality can further complicate the single pane of management enjoyed by many VMware customers today. In other words, many customers have decided that the responsibility of the server administrator ends at the NIC. Everything inside the server (up to the physical NIC) is the responsibility of the server administrator. Everything outside the server (down to the physical NIC port) is the responsibility of the network administrator. As a result, it is not uncommon to have server administrators managing the virtual network technologies contained within the server.

The needs for clean lines of responsibility are critical to large organizations, but oversimplification can introduce a lack of features and functionality. In addition, it can impact the ability to secure network access to specific servers/ports and deliver adequate quality of service to a specific port defined by a network policy.

The 1000V delivers the robust features demanded by many organizations, but does so by complicating the clean line of responsibility within the ESX host. As a result, the use of the 1000V will often require additional trained network support personnel to verify proper installation and configuration; thus overlapping the line between a server and network administrator.

Although the VMware support engineer can verify functionality from the vCenter management platform, a network engineer may need to be involved to assist in troubleshooting of the Nexus 1000V virtual switch. The network engineer may need to:

- Verify that the add-on software has been properly installed and configured on the ESX server
- Verify connectivity between the VSM and VEM
- Verify connectivity between the VSM and vCenter Server using the appropriate plug-in extension and authentication credentials using CLI commands on the VSM
- Log into the ESX server to review the log files for the Nexus 1000V module

The need to have a network engineer log into the ESX server to gather troubleshooting information dissolves the clear boundary of server versus network administrator responsibilities.

The design model of the VSM for the Nexus 1000V is as a virtual appliance. The VSM "virtual appliance" can be installed as a virtual machine on a server. A server administrator can manipulate the VSM at any time. It is conceivable that a server administrator could impact the availability of the VSM and associated "virtual line cards" by performing routine tasks on the ESX host. With this in mind, customers implementing this design must decide which team manages the integrity of VSM virtual appliance.

Summary of Virtual Switching Design Considerations with HP BladeSystem

Today, there are two main considerations when implementing a vSphere distributed virtual switch on HP BladeSystem. The customer may chose to use a VMware vNetwork Distributed Switch (vDS) or the Cisco Nexus 1000V 3rd party Distributed Switch. These considerations can have significant impact to the customer.

The drawings on page 15 and 16 articulate the differences in these implementations and their impact to the datacenter infrastructure.

Cisco Nexus 1000V 3rd-party Distributed Virtual Switch

The use of a Cisco Nexus 1000V Distributed Virtual Switch within the vSphere ESX host gives the Cisco network administrator a fully managed virtual switch. This fully managed switch works well with any HP BladeSystem Interconnect device. However, if the Network Administrator were to enable VN-Link with VNTag, the only interconnect device (shipping today) that will work with the 1000V is the HP BladeSystem pass-through module. The use of VNTag would even preclude Cisco's own CBS 3x20 switches or existing rack mounted Cisco Catalyst switches. See below for a visual depiction of this situation.

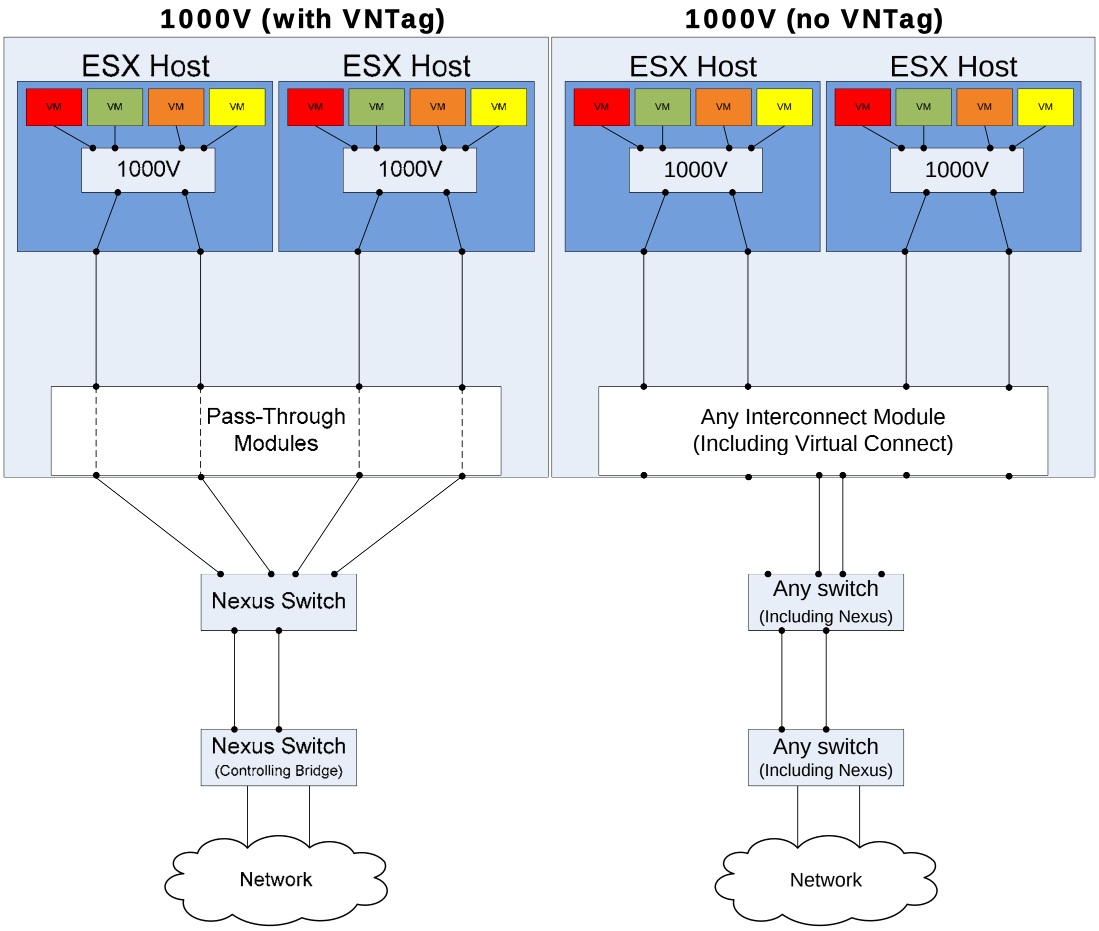
1000V (with VNTag)
1000V (no VNTag)
ESX Host
ESX Host
ESX Host
ESX Host
VM
1000V
Pass-Through Modules
Any Interconnect Module (Including Virtual Connect)
Nexus Switch
Nexus Switch (Controlling Bridge)
Any switch (Including Nexus)
Any switch (Including Nexus)
Network
Network

VMware Distributed Switch vDS

VMware's vNetwork Distributed Switch (vDS) provides many benefits over the legacy vSwitch technology in VMware ESX 3.5. The vDS still provides basic virtual machine switching capabilities, but with vSphere the vDS is now clustered across two or more ESX hosts. The configuration information is stored in the vCenter database and it is guaranteed to be consistent across all hosts in the cluster. As a result, the solution for delivering consistent uplinks to the ESX host is critical. While the VMware vNetwork Distributed Switch can plug-into any customer's existing virtual infrastructure and processes today, HP Virtual Connect Flex-10 provides a mechanism for presenting consistent "networks" and VLANs to the hosts contained in the ESX cluster.

VMware vNetwork Distributed Switch (vDS)

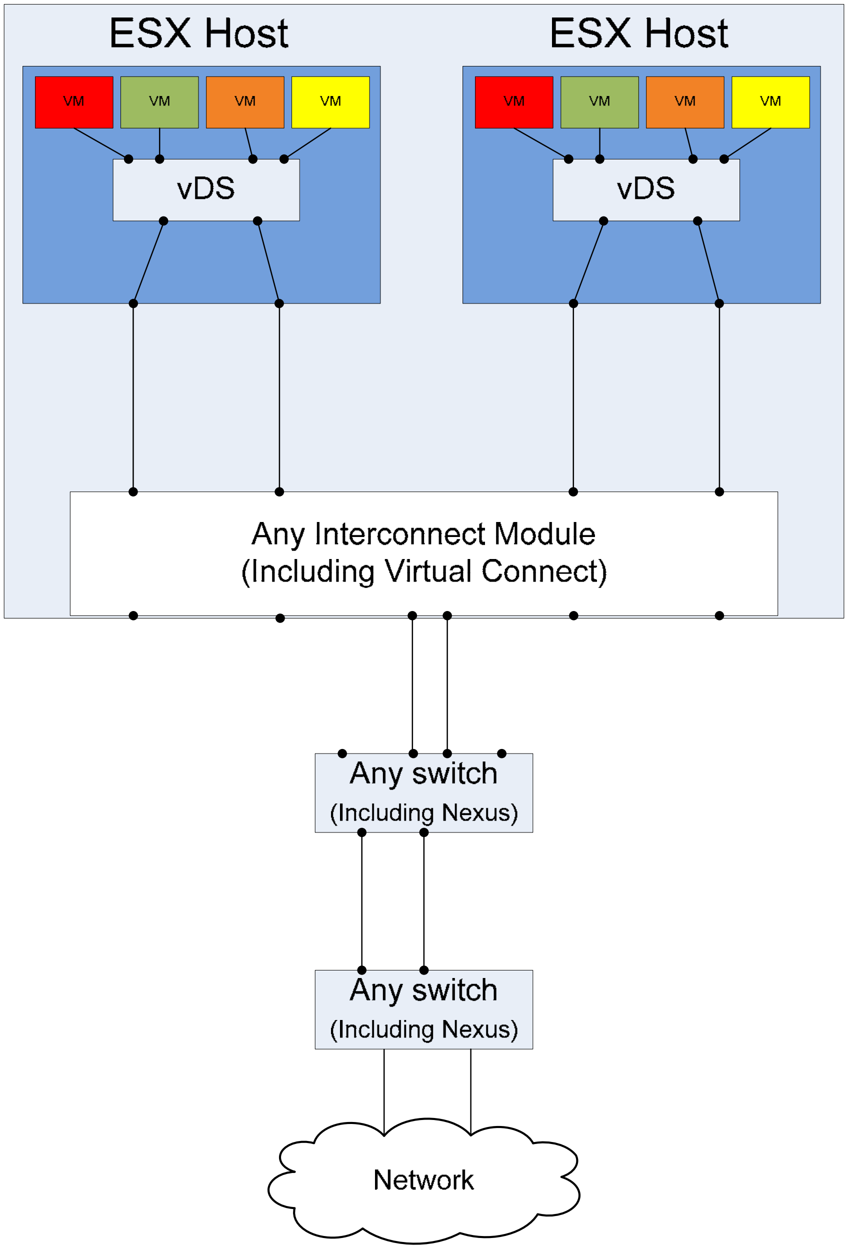

Recommendation

HP developed Virtual Connect to provide operational benefits to customers. Virtual Connect simplifies networking infrastructures by reducing the number of cables coming from the server enclosure, without adding any switches to manage. It simplifies data center management by cleanly separating the role of server management from the management of LANs and SANs; enabling Server Administrators to be virtually self-sufficient. With VC, the network administrator provisions the network VLANs to the Virtual Connect infrastructure and the server administrator connects the NICs to the appropriate network. The network administrator essentially “wires” the environment one time and the server administrator simply connects to the provided networks as needed. Virtual Connect Flex-10 extends these benefits further by offering customers more network capacity without added hardware along with adjustable network bandwidth to meet the needs of the application using it.

While there are many benefits to the Cisco Nexus 1000V, the 1000V does add some additional points of management to an ESX deployment. A 1000V running on each ESX host must be maintained and updated by a Cisco NX-OS savvy resource (usually on the network team). With this in mind, an ESX administrator would likely be required to rely on the network team to setup virtual ports for virtual machines, VLAN configurations, etc.

From a point of management and troubleshooting perspective, the 1000V adds at least one virtualized edge switch (in the form of a VEM) to every ESX server deployed. In addition, a supervisory module (VSM) running NX-OS must be maintained for up to 64 VEMs. As a result, for every ESX server deployed, a network administrator inherits at least one additional Cisco switch (as a virtual line card) to manage. By comparison, the VMware virtual distributed switch (vDS) has similar points of management, but the vDS interface is designed to be managed & maintained by an ESX administrator through vCenter where it can be managed as one switch (in the vSphere DataCenter). This single vDS management point can simplify administration overhead while performing core virtual switching functionality. This functionality is delivered at a lower overall cost and with increased virtual networking efficiency. The built-in VMware vDS does not require the engagement of the network team to complete tasks or potentially troubleshoot the 1000V software.

To summarize, here are the key considerations when choosing virtual network infrastructure components contained within the HP BladeSystem chassis:

vSphere distributed virtual switch considerations:

- The vDS is lower cost than the Cisco Nexus 1000V
- The vDS does not require a certified network expert to configure, manage, and maintain
- The vDS (in all modes of operation) will work with any HP BladeSystem interconnect device, including HP Virtual Connect
- The vDS allows a mature VMware IT shop to continue using in-place and proven network Change Management processes that does not complicate existing server Change Management processes.

In addition, the BladeSystem interconnect module considerations:

- Virtual Connect does things that traditional network devices cannot do (i.e., pre-provisioning networks, LUNs, moving server profiles)
- Virtual Connect Flex-10 saves money vs. 1Gb networking (when using 2+ NICs) and delivers 10Gb performance
- Virtual Connect is a key component of the Adaptive Infrastructure; enabling higher level capabilities
- Virtual Connect is compatible with any LAN, including Cisco Nexus
- Virtual Connect is the cornerstone for future HP technologies and continue evolve with industry standards

Summary

In summary, the use of distributed virtual switch technology over the traditional virtual switch technology is very compelling. The use of the Nexus 1000V distributed virtual switch may offer some compelling capabilities for some customers, but it may introduce challenges for others. If a customer chooses to implement the 1000V, the customer can safely do so. The 1000V is compatible with all HP interconnect devices (including Virtual Connect). If a customer chooses to implement the 1000V with VNTag, all switch (including Cisco 3x20) and Virtual Connect interconnect devices will no longer function and must be replaced with the pass-through interconnect device.

FAQ

Does HP Virtual Connect work with Cisco Nexus switches?

YES - Virtual Connect (VC) and VC Flex-10 will work with and interoperates with any industry standard Ethernet installation including the Cisco Nexus product line.

Background:
The Cisco Nexus product line is the first Cisco product to support Converged Enhanced Ethernet (CEE) protocols. The primary purpose of these protocols is to support Fibre Channel over Ethernet (FCoE). FCoE requires new "lossless" protocols to support the requirements of Fibre Channel. Cisco calls their version of these protocols DCE (Data Center Ethernet) and may add some additional proprietary features/extensions. The T11 standards body has approved the storage standards to support FCoE; however, IEEE is still working on the Ethernet standards necessary to support CEE or DCE. They are only in draft form and have not been formalized yet.

The key objective of these new protocols is to reduce the number of switches, cables, transceivers, and adapters needed to run Ethernet and Fibre Channel (FC) in a data center. All of the adapters and the majority of the cables in a data center are located between the servers and the first set of Ethernet and FC switches. Most customers can eliminate more adapters, cables, and switches by using BladeSystem with VC Flex-10 than they will ever be able to with FCoE over CEE/DCE - and they can do it today, without waiting for any new IEEE standards, inventions or testing and without replacing their current Ethernet or FC infrastructure. Future versions of Virtual Connect Flex-10 will also support CEE and offer even more savings than VC Flex-10 does today.

If a user is considering a Nexus infrastructure, then Virtual Connect Flex-10 is a good way to connect the servers to it. Customers choosing to leverage BladeSystem and Flex-10 will often find it much more cost-effective to use ProCurve, Brocade, or Cisco traditional Ethernet and Fibre Channel switches for the aggregation layer.

Does HP Virtual Connect work with Cisco Nexus 1000V?

YES - Virtual Connect (VC) interoperates with the Cisco Nexus 1000V module; provided that the configuration is utilizing industry standard Ethernet protocols.

Reference Material

HP Virtual Connect essential documents:

- HP Virtual Connect: Common Myths, Misperceptions, and Objections - http://bizsupport.austin.hp.com/bc/docs/support/SupportManual/c01598869/c01598869.pdf
- HP Virtual Connect: Ethernet Networking Scenario Cookbook: Single Domain Scenarios – http://bizsupport.austin.hp.com/bc/docs/support/SupportManual/c01471917/c01471917.pdf
- HP Virtual Connect: Users Guide - http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c01730611/c01730611.pdf
- HP Virtual Connect for the Cisco Network Administrator - http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01386629/c01386629.pdf

For more information on VMware vSphere 4, see: http://www.vmware.com/products/vsphere/.

For more information on vSphere virtual networking, see the Virtual Networking Concepts document at: http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf

For more information on troubleshooting the Cisco Nexus 1000V Troubleshooting Guide: http://www.cisco.com/en/US/docs/switches/datacenter/nexus1000/sw/4_0/troubleshooting/configuration/guide/n1000v_troubleshooting.html

Technology for better business outcomes

4AA2-9442ENW, October 2009