

HP Virtual Connect: Common Myths, Misperceptions, and Objections

“A technical discussion of common myths, misperceptions, and objections to the deployment and use of HP Virtual Connect technology.”



Table of Contents

Abstract.....	3
Target Audience	3
Prerequisites and Versioning	3
Introduction.....	4
Definitions.....	5
Myths, Misperceptions, and Objections.....	7
#1: VC Ethernet is just another switch.....	7
Comparing VC and Server Virtualization Hypervisor Networking Technology.....	7
#2: VC Ethernet is really a ProCurve switch and may not be interoperable with 3rd party network switches	9
#3: VC Ethernet doesn't support Spanning Tree (STP).....	9
#4: VC can cause duplicate MACs and WWNs on the network.....	9
#5: VC users cannot leverage existing network management tools (for example, CiscoWorks).....	9
#6: VC doesn't provide secure external management.....	10
#7: VC Ethernet doesn't support Private VLANs	10
#8: VC doesn't provide deterministic load balancing for multiple LACP channels.....	10
#9: VC Ethernet doesn't support VLAN Trunking to Server Blade NICs	10
#10: VC Ethernet and Fibre Channel modules require a clunky hardware failure recovery (RMA) process	11
#11: VC doesn't provide network visibility into the VC Domain for network administrators	11
#12: Many customers experience problems with VC deployment consistency.....	11
#13: VC doesn't support non-disruptive firmware upgrades ("hot code load").....	12
#14: HP server blade NICs stay active even after VC Ethernet uplink failure	12
#15: VC is proprietary to HP; VC locks a user to a single blade vendor	12
#16: VC modules form a master/slave relationship with separate IP addresses that aren't transferred during a failover. When multiple VC modules are present, only one module is elected to forward traffic ..	13
#17: VC only supports a maximum of 64 VLANs.....	13
#18: VC Ethernet doesn't provide Layer 3 routing capabilities	13
#19: VC's CX4 cables are not ideal for 10Gbe uplinks because of distance limitation.....	14
#20: VC doesn't support stacking multiple VC modules	14
#21: VC doesn't offer Access Control Lists or VLAN Access Control Lists (ACLs or VACLs)	14
#22: VC Ethernet doesn't support user configurable Quality of Service (QoS) features	14

#23: VC Ethernet doesn't provide diagnostic tools (SPAN).....	15
#24: VC Ethernet doesn't support the Cisco Discovery Protocol (CDP)	15
#25: VC requires a separate web management window to manage each VC module	15
#26: VC doesn't support IGMP Snooping	15
#27: VC Fibre Channel doesn't support nested NPIV	15
#28: Virtualized MAC and WWN (NATing) features on some switches (for example, Cisco's FlexAttach* feature) provide the same benefits as VC's Managed MAC addresses and WWNs	16
#29: DHCP Option 82 provides the same server redundancy features as Virtual Connect	17
#30: VC-FC doesn't provide login distribution and failover on FC uplinks to the SAN	19
#31: All VC-FC uplinks have to be connected to the same SAN fabric.....	19
#32: VC implements an immature loop avoidance mechanism	19
#33: VC Uplink failures require re-convergence on the external network and may cause dropped server sessions	19
#34: Cisco's N-Port Virtualization (NPV)* or Brocade Access Gateway** provide all the same advantages as VC-FC	20
#35: Cisco VFrame Data Center provides the same capabilities as VC	20
#36: VC Ethernet can't be connected to Cisco 6500 switches using Virtual Switching System (VSS)	20
#37: VC Ethernet does not support Unidirectional Link Detection (UDLD).....	21
#38: VC Ethernet only provides port counters on uplinks	21
Unique Features Provided By HP Virtual Connect	22
Managed Server Identities	22
Internal Server Identity	22
External Server Identity.....	22
Preprovisioning Using Managed Server Identities	22
"LAN Safe" Network Connectivity.....	23
Server Adds, Moves, and Replacements are Transparent to LAN & SAN	23
Summary of the Virtual Connect Capabilities	24
Additional Resources and References	27
About the Author	27
Appendixes	28
Appendix A: Frequently Asked Questions	28

Abstract

This whitepaper discusses some of the common myths, misperceptions, and objections to the deployment and use of HP's Virtual Connect technology in data center networks. Technical answers are provided for these common assertions, whether correct or incorrect, to help the reader sort the facts. In addition, this paper includes an in-depth discussion and comparison of the Virtual Connect feature set and capabilities when compared to the deployment of traditional LAN and SAN switches.

Target Audience

The target audiences of this whitepaper are current Virtual Connect users who would like to learn more about the capabilities of Virtual Connect and potential users who are evaluating and testing Virtual Connect for possible adoption. This paper is also targeted for any audience who may have received incorrect information regarding the features, functions, and capabilities of the Virtual Connect technology suite.

Prerequisites and Versioning

It is assumed that the reader is already familiar with Ethernet networking terminology, features and device operation and that the reader is familiar with the basics of HP BladeSystem c-Class enclosures, HP BladeSystem c-Class server blades, and HP BladeSystem Virtual Connect. For additional information on these HP BladeSystem c-Class components, please visit:

<http://www.hp.com/go/bladesystem> &
<http://h18004.www1.hp.com/products/blades/components/c-class-interconnects.html>

Recommended Prerequisite Reading:

- [Non-technical Summary of Virtual Connect Technology](#)
- [White Paper: How to implement Virtual Connect](#)
- [Virtual Connect for the Cisco Network Administrator](#)
- [Latest Virtual Connect Documentation](#) (see User Guide)

This whitepaper was written based on the features provided in Virtual Connect Ethernet firmware version 1.3x and Virtual Connect Fibre Channel firmware version 1.2x. Newer releases of firmware may introduce new features or may introduce changes to the way existing features work. For any discrepancies between the information in this paper and actual device operation, HP recommends the Administrator refer to the Virtual Connect User Guide and Release Notes matching the firmware version in use. Virtual Connect firmware documentation can be found under the "Install Your Solution" tab at www.hp.com/go/bladesystem/documentation.

Introduction

Virtual Connect is an innovative server identity virtualization and I/O management product for HP BladeSystem c-Class customers. It first shipped in February 2007. Virtual Connect has been deployed successfully in thousands of customer environments, many with large and growing installations. It is one of the key reasons HP BladeSystem represents approximately half of the blades market today. Virtual Connect was developed as a direct result of customer requests for a better way to manage server blade network and storage connectivity and virtualization. As a result, Virtual Connect addresses several key challenges by providing the following capabilities and features:

- 1) Virtual Connect reduces the number of cables required to connect servers to LANs and SANs without having to increase the number of traditional Ethernet or Fibre Channel switches to manage.
 - a. Presents the entire c-Class enclosure to the LAN and SAN in the same way as a server virtualization hypervisor (for example, VMware, ESX, Microsoft HyperV, Citrix Xen, etc.)
 - b. Doesn't require configuring all the traditional switch-to-switch protocols (STP, VTP, etc.)
- 2) Virtual Connect enables a clean separation between the server infrastructure and the network and storage infrastructure.
 - a. Enables a server administrator to become self-sufficient when making adds, moves, and changes of servers since no corresponding changes to Ethernet switch or Fibre Channel switch configurations are required.
 - b. Reduces distractions for network and storage administrators who are no longer repeatedly interrupted with server change requests.
- 3) Virtual Connect manages a server's internal identity and a server's external identity to enable transparent server adds, moves, and replacements.
 - a. Internal Identity: Server hardware replacement is transparent to the OS because Virtual Connect virtualizes, manages, and maintains constant, the server's internal identity, which includes the server's serial number, UUID, BIOS settings, and FC Boot parameters.
 - b. External Identity: Server adds, moves, and replacements in the data center are transparent to external LANs and SANs because Virtual Connect manages, and maintains constant, the server's external identity, which includes the server's Ethernet MAC addresses, Fibre Channel WWNs, LAN assignments (VLANs), and SAN assignments (fabrics).

As with any new product that introduces a new way of solving old problems, many assumptions are often made regarding Virtual Connect's operation without fully understanding how VC's technology works. Fundamentally, Virtual Connect presents itself and the c-Class enclosure to the LAN and SAN in a manner very similar to the way a hypervisor presents itself to the network. The key difference is that Virtual Connect provides this capability as a hardware solution rather than a software layer in the server.

Virtual Connect allows multiple hosts to share a common Ethernet path to the LAN and a common Fibre Channel path to the SAN without having to configure and deploy lots of traditional Ethernet and Fibre Channel switches. This technique has been proven out with many customers over the years. Virtual Connect has been successfully installed in and used with a variety of different vendor data center networks, such as Cisco and Brocade networks. HP provides extensive installation and integration documentation and professional services to enable successful Virtual Connect deployment.

Many potential implementers receive incorrect information regarding Virtual Connect's capabilities. This is generally due to one of three things:

- 1) A lack of understanding that Virtual Connect is part of the server infrastructure (like server virtualization hypervisor)
- 2) A general lack of understanding of how the Virtual Connect technology works and the specific features and capabilities it provides
- 3) A competitor's occasional meritless, erroneous, and unsubstantiated attack on Virtual Connect's features and capabilities

This paper provides factual responses to many of these points to provide the VC implementer with a technically accurate description of VC usage and capabilities. This paper is intended to supplement other HP documentation.

Definitions

ACL (VACL)	Access Control List or VLAN Access Control List: A set of rules that allows or disallows network traffic to flow between network devices
BPDU	Bridge Protocol Data Unit: A spanning tree configuration frame exchanged between switches in the same spanning tree domain
CDP	Cisco Discovery Protocol: A proprietary Cisco protocol used to exchange neighbor information between two directly connected Cisco devices
CX-4	An industry standard cabling specification used by VC for network connectivity using 10 Gbit Ethernet over copper
Downlink	An internal port (enclosure midplane) on a blade switch or Virtual Connect Ethernet module that directly connects to a server blade's NIC port
External Network	The network and associated network devices external to the VC domain
Hypervisor	A virtual machine hypervisor such as VMware ESX, Microsoft Hyper-V, Citrix Xen, etc.
Internal cross-connect	A non-visible port that interconnects two horizontally adjacent VC-Enet modules
LACP	Link Aggregation Control Protocol: An 802.3ad Link Aggregation configuration frame exchanged between two devices that form a port trunk/channel between them
LAG	Link Aggregation Group. 802.3ad terminology for a port trunk/channel group
LLDP	Link Layer Discovery Protocol. An IEEE 802.1ab protocol that provides CDP-like functionality
Logical Path	A single physical port or a single port channel. Both represent a single communication path.
LOM	LAN on Motherboard. A NIC embedded on the system board of a server.
NAT	Network Address Translation: A feature that allows a network device (such as a switch or router) to replace/rewrite addresses within a frame with a different address
NPIV	N-Port ID Virtualization: ANSI T11 feature for Fibre Channel that allows multiple WWNs to login to the fabric over a single N-Port.
Port Trunk (channel group)	A group of two or more ports that operate as a single logical port and single logical path for the purposes of load balancing. 802.3ad and EtherChannel are both port trunking technologies
Quality of Service (QoS)	A very broad term associated with network traffic classification, prioritization, queuing, marking, etc
Server Profile	An object within the Virtual Connect domain that is assigned to a server bay and contains the server's LAN and SAN connectivity settings (vNet assignments, managed MAC addresses & WWNs, server boot parameters, PXE configuration, and Fibre Channel boot parameters).
SFP	A hot-pluggable modular 1 GbE port. Pluggable modules allow for electrical or optical connectivity at 1 Gbit speeds
Shared Uplink Set (SUS)	The term used by Virtual Connect to configure one or more VC uplinks as a VLAN trunk connected to a switch employing IEEE 802.1Q VLAN trunking
Stacking Link	A link that directly connects two VC Ethernet ports from different VC Ethernet

	modules that belong to the same VC domain
Uplink	A external faceplate port on a blade switch or Virtual Connect Ethernet module that directly connects to an external network device
UUID	Universally Unique Identifier: a 128 bit globally unique object identifier that is used by things like Operating Systems and applications to identify specific and individual computing devices (servers)
VC	Virtual Connect: Broad term used to reference all the Virtual Connect components as a whole – Ethernet & Fibre Channel modules and Virtual Connect Manager.
VC-Enet	A Virtual Connect Ethernet module
VC-FC	A Virtual Connect Fibre Channel module
VCM	Virtual Connect Manager: The user interface, web or CLI, used to manage a Virtual Connect domain
Virtual Connect Domain	All VC Fibre Channel modules and all stacked VC-Enet modules within the same enclosure and under the control of the same Virtual Connect Manager
Virtual Connect Network (vNet)	A logical grouping of VC ports (downlinks or downlinks & uplinks) that comprise a single layer 2 network or broadcast domain.
VC Downlink	Non-visible ports that are directly connected to server NIC ports through the enclosure midplane.
VCEM	Virtual Connect Enterprise Manager: A separate software product that extends management to as many as 100 VC domains from a single console.
VC Uplink	Visible ports on the VC-Enet module faceplate that provide external connectivity for the server blades inside the enclosure.
VLAN Trunk	A single physical port or a single port channel with VLAN tagging enabled. Used to provide connectivity to one or more VLANs over the same logical path.
vSwitch	A hypervisor virtual switch. A software implementation of a layer 2 bridge as used by virtual server hypervisors
WWN	World Wide Name: Equivalent to a MAC address for a Fibre Channel device.
XFP	A hot-pluggable modular 10 GbE port. Pluggable modules allow for electrical or optical connectivity at 10 Gbit speeds

Myths, Misperceptions, and Objections

#1: VC Ethernet is just another switch

Incorrect: While VC uses tried-and-true, IEEE standard, Layer 2 bridging functionality, its primary purpose is to provide many server virtualization and management features that are non-existent in traditional switches. VC may perform some functions like a traditional switch, however, VC has many additional features which clearly distinguish it from a traditional switch. Likewise, server virtualization hypervisors (for example, VMware ESX, Microsoft Hyper-V, Citrix Xen) perform some functions of a traditional switch but, like VC, have many additional features which clearly distinguish them from a traditional switch. As a result, it is incorrect to say that either technology, VC or hypervisor virtual switching, is “just another switch”.

VC and server virtualization hypervisors are very similar in the networking functionality that they provide to servers; a hypervisor provides it for virtual servers and VC provides it for physical HP server blades. In the same way that a hypervisor provides this functionality in a way that interoperates with the external network, VC also provides this interoperable connectivity between HP server blades and the external network. Virtual Connect is not called a “switch” because it is not configured, deployed, or managed as a switch and does not present itself to the external network as a switch – again, much like a hypervisor. When Virtual Connect is linked to the external network, the external network “sees” the same behavior from VC as it “sees” when a server hosting a hypervisor is connected to the external network. Since VC is not configured, deployed, or managed like a traditional switch and presents itself to the network as an endpoint (like a server), it is incorrect to call VC a “switch”. For a more thorough discussion of the similarities of VC and hypervisors from a networking technology perspective, see the following section – “Comparing VC and Server Virtualization Hypervisor Networking Technology”.

NOTE: Even though VC and server virtualization hypervisor networking technology similarities are being discussed, the two products provide a solution for completely different problems in the data center. As such, VC and server virtualization hypervisors (for example, VMware ESX, Microsoft Hyper-V, Citrix XEN, etc.) work together very well to provide a robust solution.

Comparing VC and Server Virtualization Hypervisor Networking Technology

One method of understanding how Virtual Connect operates on the LAN is to compare the Virtual Connect networking components and their functionality to the networking components of a server virtualization hypervisor (for example, VMware ESX, Microsoft Hyper-V, Citrix Xen). Since the networking technology behind hypervisors is commonly understood and accepted by many customers, understanding the many similarities between VC and hypervisors will help an implementer have a better understanding of how Virtual Connect looks to, and behaves on, the external network. Just to be clear, Virtual Connect and hypervisors are fundamentally different products and address completely different needs within the data center. This comparison is strictly about understanding the similarities between the two products in regards to networking technology in order to better understand Virtual Connect.

A Description of the Hypervisor Components:


Referencing figure 1 below, the hypervisor host (left) is a single physical server running a server virtualization hypervisor (for example, VMware ESX, Microsoft Hyper-V, Citrix Xen) that allows the physical server to host one or more instances of a virtual server, called a Virtual Machine (VM). In addition, the hypervisor host provides external network connectivity to the internal servers (VMs) using a virtual (software) implementation of a layer 2 bridge, called a vSwitch. The VM virtual NICs (vNics) are assigned to one of the vSwitches and the vSwitches are then associated with real physical NICs residing in I/O slots on the hypervisor host. The vSwitches can have one or more physical NICs (uplinks) assigned to them to provide external network connectivity. If more than one physical NIC is

assigned to the same vSwitch, network redundancy and/or load balancing is provided for the internal servers (VMs) assigned to that vSwitch. The physical NICs then present one or more MAC addresses to the external network, depending on the number of VMs communicating to the external network through each physical NIC.

A Comparative Description of the VC Components:

Referencing figure 1 below, the c-Class enclosure (right) is a single physical enclosure that hosts one or more real physical servers, called a server blade. In addition, the c-Class enclosure provides external network connectivity to the internal servers (server blades) using a hardware implementation of a layer 2 bridge, called a Virtual Connect Ethernet network (vNet). The server blade's physical NICs (pNics) are assigned to one of the vNets and the vNets are then associated with real physical VC uplink ports from VC-Enet modules residing in the I/O bays on the c-Class enclosure. The vNets can have one or more VC uplinks assigned to them to provide external network connectivity. If more than one VC uplink is assigned to the same vNet, network redundancy and/or load balancing is provided for the internal servers (server blades) assigned to that vNet. The VC uplinks then present one or more MAC addresses to the external network, depending on the number of server blades communicating to the external network through each VC uplink.

Figure 1. Hypervisor Networking Technology Compared to Virtual Connect Enclosure



After comparing the components and their functionality, it is obvious why many customers treat a c-Class enclosure with Virtual Connect the same way they would a single host running a hypervisor. In other words, VC allows an entire c-Class blade enclosure to look to the network like one big hypervisor host. From a network redundancy and load balancing perspective, from a security perspective, and from a port monitoring perspective, VC simplifies the network connectivity for an entire c-Class blade enclosure and makes it behave on the network like a single host running a hypervisor.

#2: VC Ethernet is really a ProCurve switch and may not be interoperable with 3rd party network switches

Incorrect: VC is not a ProCurve switch product and VC is interoperable with any other IEEE compliant network device. VC is a product engineered, developed and sold by the HP BladeSystem division independently of the HP ProCurve division. VC operates according to IEEE standards and, like a ProCurve switch, will interoperate with any networking device that is also IEEE compliant. Virtual Connect provides network-vendor-independent connectivity for an HP BladeSystem c-Class enclosure.

#3: VC Ethernet doesn't support Spanning Tree (STP)

Correct: Much to the delight of VC users, Spanning Tree support on VC is not needed. VC provides HP server blade network connectivity just like a hypervisor provides virtual server network connectivity and neither of these technologies require Spanning Tree support. VC doesn't have to support Spanning Tree just like hypervisor hosts don't have to support it, yet both provide network redundancy and load balancing. Just like a hypervisor host, VC provides network redundancy and load balancing features that are modeled after NIC Teaming/bonding technology instead of switch technologies like Spanning Tree with error-prone configurations. A Spanning Tree configuration error on any single switch in the data center can negatively affect any other connected switch in the network, in addition to all servers connected to the same network. With Virtual Connect, any redundancy and load balancing configuration problems only affect a single blade enclosure*.

Fundamentally, VC doesn't require support for protocols like STP because VC presents itself to the network as a "termination endpoint", as does a typical server or a hypervisor host. VC is not and does not present itself as a "transit device", as does a traditional switch.

(See response to question 32 for a discussion on loop avoidance.)

* Up to four enclosures if using VC enclosure stacking.

#4: VC can cause duplicate MACs and WWNs on the network.

Incorrect: Virtual Connect Manager (VCM) prevents duplicate MAC addresses and WWNs on the network for servers in the same VC Domain and Virtual Connect Enterprise Manager (VCEM) prevents duplicate MAC addresses and WWNs for servers across multiple VC Domains. Within a VC Domain, all MACs and WWNs are restricted to a single server port at any one time. Regardless of how physical servers are inserted, removed, swapped, or replaced, Virtual Connect prevents the same Virtual Connect Managed MAC address or WWN from being used on more than a single NIC or HBA port. A user could introduce duplicate MACs and WWNs on the network by improperly selecting an address range already in use by another VC domain. To ensure that this problem does not occur, customers have the option of using Virtual Connect Enterprise Manager (VCEM) to manage up to 100 VC Domains within the data center.

In addition ensuring that all VC managed MACs and WWNs are unique, VCEM also provides automated server recovery and server movement across multiple VC domains.

#5: VC users cannot leverage existing network management tools (for example, CiscoWorks)

Incorrect: Virtual Connect supports configuration scripting via a CLI interface (SSH) and Virtual Connect supports monitoring using SNMP. Any management tools that support CLI scripting can be

used to remotely configure Virtual Connect. Any management tool that supports SNMP can be used to monitor Virtual Connect.

VC Ethernet supports applicable groups within the following MIBs: Compaq Host MIB, Compaq System Info MIB, RFC 3418 SNMPv2-MIB, RFC 2863 IF-MIB, RFC 4188 BRIDGE-MIB. In addition to local statistics and SNMP polling of statistics, VC provides SNMP traps for events that cause VC Domain status changes. For additional details on the support MIBs and traps, please refer to the Virtual Connect User Guide for the appropriate firmware version in use.

#6: VC doesn't provide secure external management

Incorrect: Virtual Connect provides SSH (CLI) and SSL (Web GUI) support for secure remote management. VC also supports LDAP (for example, Microsoft Active Directory and OpenLDAP) for centralized user management.

#7: VC Ethernet doesn't support Private VLANs

Incorrect: As of VC firmware version 1.31, Virtual Connect provides Private VLAN support. The VC feature is called "Private Networks". See the VC firmware version 1.31 User Guide for more details.

In addition to VC's own support for Private Networks (Private VLANs), VC can also be configured to extend the Private VLAN configuration from external switches. For additional information, see the section entitled "Private VLANs" on page 42 of the whitepaper "[Virtual Connect for the Cisco Network Administrator](#)".

#8: VC doesn't provide deterministic load balancing for multiple LACP channels

Incorrect: Virtual Connect utilizes a deterministic load balancing algorithm for frames load balanced across a single LACP channel and VC provides a deterministic algorithm for determining active vs. standby LACP channels assigned to the same Virtual Connect network (vNET).

VC's algorithm for frame load balancing within a LAG (port trunk/channel) is automatic based on the protocol information in the frame. If the frame contains Layer 4 information (for example, TCP, UDP), Virtual Connect will use it and Layer 3 information (source and destination IP addresses) to determine conversation streams and statistically load balance individual streams on different ports in the LAG. If a frame only contains Layer 3 information (IP addresses), Virtual Connect will use the source and destination IP addresses to determine the conversations to load balance. For all Layer 2 only frames, VC simply uses the Source and Destination MAC addresses to determine conversations to load balance across the different ports in the LAG.

When multiple LAGs are configured in a single vNet, VC determines the active LAG based on LAG bandwidth. The LAG with the most bandwidth (port speed + number of active ports) becomes the active LAG. All other LAGs in the vNet are put in standby mode (like NIC Teaming). If all LAGs are equal, VC Enet module ID (MAC address) and uplink port numbers are used to break the tie.

#9: VC Ethernet doesn't support VLAN Trunking to Server Blade NICs

Incorrect: Virtual Connect does support VLAN Trunking to Blades. VC firmware release 1.31 and above provides full support for VLAN Trunking to HP server blade NICs. Using Cisco switch port

mode descriptions, VC supports “access” mode, “trunk” mode, and “dot1qtunnel” mode to any server blade NIC. The VC administrator can choose which mode to use to customize the VC configuration for their environment

Further, the VLAN Trunking enhancements included in version 1.31 actually provides more flexibility than a traditional switch in order to provide enhanced server virtualization and transparent movement within the data center or across data centers (disaster recovery). This feature, called Mapped VLAN IDs, allows the administrator to translate tagged VLAN IDs originated by a server to the correct VLAN ID used by the external network. This VLAN translation (mapping) is completely transparent to the server blade and the external network. Mapped VLANs are controlled and configured by users with LAN administrator rights within Virtual Connect for security purposes.

#10: VC Ethernet and Fibre Channel modules require a clunky hardware failure recovery (RMA) process

Incorrect: Virtual Connect provides high availability and fault recovery using configuration check pointing/synchronization across adjacent VC Ethernet modules within each c-Class enclosure. In the unlikely case a VC module fails (Ethernet or Fibre Channel), the complete configuration is retained by the VC Domain using modules in interconnect bays 1 and 2. When the failed module is replaced, the configuration is automatically restored to the newly inserted module. In other words, VC supports plug-n-play of replacement VC modules or additional modules to expand the VC Domain’s capabilities. Because HP server blades are connected to more than one redundant VC module, “no single point of failure” configurations are easy to deploy.

VC also supports exporting the VC Domain configuration for manual configuration restoration.

#11: VC doesn’t provide network visibility into the VC Domain for network administrators

Incorrect: Virtual Connect provides several user interfaces options and features for managing and monitoring Virtual Connect to fit with the variety of methods our customers use. VC supports both a Web interface (HTTPS) and a CLI interface (SSH). In addition, VC supports per-interface statistics for every server NIC port, server HBA port, VC Ethernet uplink port, and VC Fibre Channel uplink port. These statistics can be monitored via the management interfaces or via SNMP/SMI-S polling. For example, VC supports applicable groups within the following MIBs: Compaq Host MIB, Compaq System Info MIB, RFC 3418 SNMPv2-MIB, RFC 2863 IF-MIB, RFC 4188 BRIDGE-MIB. In addition to local statistics and SNMP polling of statistics, VC provides SNMP traps for events that cause VC Domain status changes.

Virtual Connect also supports port mirroring, to an external network analyzer, of Ethernet traffic to/from any server NIC port(s).

#12: Many customers experience problems with VC deployment consistency

Incorrect: Any mature, flexible, and feature-rich product provides the administrator with options for configuring and customizing it. These configurations and customizations should be tested and methodically applied to the product as it is deployed by the user. Virtual Connect is just such a product and HP always recommends that an administrator purposefully customize and deploy a VC configuration that is tailored to their environment.

To help simplify and ensure configuration consistency across similarly configured VC Domains, HP provides enhanced configuration features. Virtual Connect allows VC Domain configurations to be exported from a configured VC Domain and then imported on a non-configured VC Domain. Scripting via the VC CLI can also be used to deploy consistent VC Domain configurations across multiple enclosures. Virtual Connect Enterprise Manager also provides enhanced configuration consistency across VC Domains that are grouped together as “like” configurations.

#13: VC doesn't support non-disruptive firmware upgrades (“hot code load”)

Incorrect: Virtual Connect allows the administrator to upload firmware to each individual VC Ethernet or Fibre Channel module without interfering with that module's operation or the operation of any other VC module. Once the new firmware is uploaded, VC allows the administrator to choose when to activate the new firmware on a module-by-module basis (usually during a change window). If using a “no single point of failure” configuration for the HP server blades, individual VC modules may have their firmware activated while other VC modules maintain connectivity for HP server blades.

When upgrading VC firmware, there are two VC components to consider – the VC Manager interface (analogous to a supervisor module on a Cisco 6500) and the individual VC modules (analogous to the individual switch modules inserted into a Cisco 6500 chassis). The VC Ethernet modules in bay 1 and bay 2 of the c-Class enclosure work together to provide redundancy for the VC Manager interface (like redundant supervisor modules in a Cisco 6500). A customer can choose to upgrade VC Ethernet in bay 1 while VC Ethernet in bay 2 runs as the active VC Manager. This provides a means for upgrading VC firmware while maintaining an active VC Manager. In addition, all server blades are connected to multiple VC Ethernet and VC Fibre Channel modules when using HP's best practices. Since server NICs and HBAs are redundantly connected to VC Ethernet and Fibre Channel modules, each VC module can be individually upgraded and activated while adjacent VC modules in the enclosure provide active connectivity for the server (when using NIC Teaming for Ethernet and MPIO for Fibre Channel). This is analogous to a server having NICs connected to multiple switch modules in a Cisco 6500 and independently upgrading each of the switch modules.

#14: HP server blade NICs stay active even after VC Ethernet uplink failure

Incorrect: Virtual Connect provides many features for ensuring highly available network connectivity for HP server blades. One feature, SmartLink, is used to disable a server blade NIC port anytime the NIC is connected to an external network where all VC uplink(s) have failed. In other words, VC can be configured to proactively disable a server NIC port whenever the server NIC is isolated from the external network. VC's SmartLink feature, combined with NIC Teaming on the server, allows for highly available network configuration with no single point of failure.

#15: VC is proprietary to HP; VC locks a user to a single blade vendor

Correct; Incorrect: Virtual Connect is one of many HP products that provide HP customers with patented features and functionality that no other server blade or networking vendor offers. Customers that desire these enhanced features and products gladly deploy HP products. In today's market, IT managers understand that IT efficiency and manageability are critical to their success. As a result, most customers do not choose to deploy features based on the “lowest common denominator” method. As an example, many customers choose to deploy equipment from a single network equipment vendor within the data center based on the feature set provided even when those features

are not provided by any other network equipment vendor. Choosing a single server blade vendor, such as HP, based on their providing valuable and unique feature sets is no different.

Virtual Connect interoperates with IEEE and ANSI T11 standards and doesn't require HP proprietary devices outside of the HP Blade enclosure. VC is an optional component of HP BladeSystem and a user can choose to use other vendor products (for example, Cisco, Brocade, etc.) in place of Virtual Connect Ethernet and Fibre Channel modules. Whether using Virtual Connect or not, the same data center LANs and SANs can be used to provide network and SAN connectivity to HP server blades.

In addition, any LAN or SAN switch available for any vendor's blade chassis is unique, and therefore proprietary to that vendor's chassis. For instance, the Cisco Catalyst 3120G* only works in the HP BladeSystem c-Class family of enclosures. The Cisco Catalyst 3130G* only works in the DELL PowerEdge M1000e blade chassis. While they provide similar functionality, they are not interchangeable. Further, the Cisco Catalyst 31x0G series, available in their blade vendor specific forms, is a design unique and proprietary to Cisco.

* <http://www.cisco.com/en/US/products/ps8742/index.html>

** <http://www.cisco.com/en/US/products/ps6748/index.html>

#16: VC modules form a master/slave relationship with separate IP addresses that aren't transferred during a failover. When multiple VC modules are present, only one module is elected to forward traffic

Incorrect: Virtual Connect provides high availability for the VC domain management interface (Web or CLI) by using cluster-like technology between VC Ethernet modules. While one VC module is the "active" VC Domain management interface, the adjacent VC module is in standby mode. If the active VC module fails or is removed, the standby VC module becomes the active/master module for the VC Domain management interface. Virtual Connect can be configured so that the same IP address is always used by the active VC module within a VC Domain. This optional feature is called the "Virtual Connect Domain IP Address".

While only one VC module provides the active VC Domain management interface (Virtual Connect Manager), **all** VC Ethernet and Fibre Channel modules can be used simultaneously to provide network connectivity for HP server blades. In other words, even though only a single VC Ethernet module is elected to service the VC Manager interface, all VC Ethernet and Fibre Channel modules and all ports on these modules can be configured to forward traffic simultaneously.

#17: VC only supports a maximum of 64 VLANs

Correct: The supported limit is 64 VLANs per c-Class enclosure. The VC architecture supports 1000+ VLANs per c-Class enclosure. To date, no VC customer has requested more than 64 VLANs per c-Class enclosure. HP can increase this "supported limit" anytime in the future based on customer demand.

#18: VC Ethernet doesn't provide Layer 3 routing capabilities

Correct: Virtual Connect is not a router, therefore, Virtual Connect does not provide Layer 3 capabilities (routing). Customers that desire to route between HP server blades in the same enclosure or between HP server blades and an external device will utilize routers in their core network. Alternatively, should a customer prioritize internal routing of blade-to-blade traffic within the enclosure above the server management and virtualization features provided by Virtual Connect, HP would

recommend the customer deploy using the Cisco 3120 blade switch (with purchasable IOS upgrade for full L3 routing) or using one of the other HP c-Class GbE2c blade switch options providing Layer 3 routing capabilities (for example, HP GbE2c Layer 2/3 Ethernet Blade Switch).

Cisco's Data Center Infrastructure 2.5 Design Guide ([link](#): see page 79) advises that data center access switches and access devices (for example, Virtual Connect Ethernet) are usually deployed in Layer 2 mode (no routing). See reference above for a complete discussion.

#19: VC's CX4 cables are not ideal for 10Gbe uplinks because of distance limitation

Correct: CX4 ports have a 15 meter limitation per the IEEE specification, including those used by VC. For this reason, Virtual Connect also provides IEEE compliant fiber 10GbE ports (10GbE SR or LR) in the 1/10Gb-F Virtual Connect Ethernet module which support a distance limitation of up to 10 kilometers.

The Cisco Catalyst 3120X CX4 ports have the same 15 meter limitation. The Cisco Catalyst 3120X fiber ports offer increased distance limitation similar to VC fiber ports*.

* http://www.cisco.com/en/US/prod/collateral/switches/ps6746/ps8742/ps8749/data_sheet_c78-439133.html

#20: VC doesn't support stacking multiple VC modules

Incorrect: Virtual Connect fully supports stacking up to 8 VC Ethernet modules today and VC firmware version 1.31 and above provides support in beta for stacking up to 16 VC Ethernet modules and 16 VC Fibre Channel modules across four c-Class enclosures – all 32 managed using a single GUI or CLI interface. Any external VC Ethernet port can be used as a stacking link, an uplink to the external network, or a network analyzer port.

#21: VC doesn't offer Access Control Lists or VLAN Access Control Lists (ACLs or VACLs)

Correct: Like virtual server hypervisor vSwitches, VC isn't a traditional switch, and therefore does not currently support Access Control Lists (ACLs) or VLAN ACLs (VACLs). Depending on customer demand, the Virtual Connect architecture could support the implementation of ACLs or VACLs in a future firmware update.

As an alternative, however, VC can be configured to extend the ACL or VACL configuration from external switches. For additional information, see the section entitled "ACLs and VLAN ACLs" on page 40 of the whitepaper "[Virtual Connect for the Cisco Network Administrator](#)".

#22: VC Ethernet doesn't support user configurable Quality of Service (QoS) features

Correct: VC does not currently support any user configurable Quality of Service features. Virtual Connect uses a FIFO queuing mechanism with head of line blocking prevention. The Virtual Connect architecture supports the implementation of advanced QoS mechanisms and these features are on the roadmap for future implementation.

For users concerned about dedicated bandwidth for certain servers, HP would recommend dedicating the use of one or more of the many VC uplinks provided or HP would recommend enabling Quality of Service features on the first upstream switch port. For examples of deploying customized VC configurations to provide dedicated bandwidth to one or more servers, see figure 11 on page 28 of the whitepaper "[Virtual Connect for the Cisco Network Administrator](#)". In this figure, server blades 1, 2 and 3 are all connected to VLAN 1. However, server blade 3 has a pair of dedicated uplinks to VLAN 1 which provides dedicated bandwidth only to server blade 3. In addition, the network administrator could configure marking and traffic prioritization on the switch ports connected directly to Virtual Connect uplink ports.

#23: VC Ethernet doesn't provide diagnostic tools (SPAN)

Incorrect: VC supports port mirroring or monitoring of server NIC traffic to a VC uplink on the same VC module (equivalent to Cisco's SPAN) and VC supports port mirroring of server NIC traffic to any VC uplink on any VC module in the VC Domain (equivalent to Cisco's RSPAN). A user can also configure RSPAN on an external Cisco switch port, in conjunction with VC's port monitoring feature, to send mirrored server blade traffic to any remote device within the network.

#24: VC Ethernet doesn't support the Cisco Discovery Protocol (CDP)

Correct: VC supports Link Layer Discovery Protocol (LLDP) – the industry standard (IEEE) version of the Cisco proprietary protocol CDP. Many Cisco devices support both CDP and LLDP (for example, Cisco 3120*). The use of the IEEE standard version, LLDP, is recommended by HP to ensure customers are not locked into a proprietary protocol.

*

http://www.cisco.com/en/US/docs/switches/blades/3120/software/release/12.2_46_se/release/configuration/guide/swlldp.html

#25: VC requires a separate web management window to manage each VC module

Incorrect: A single Web management window or CLI prompt is used to manage all VC modules in the same VC Domain. If using Virtual Connect Enterprise Manager (VCEM), a single VCEM User Interface can be used to manage server profile management (creation, deletion, movement, etc.) across up to 100 VC Domains.

#26: VC doesn't support IGMP Snooping

Incorrect: VC supports IGMP snooping for IGMP versions 1 and 2.

#27: VC Fibre Channel doesn't support nested NPIV

Incorrect: This feature is provided in VC version 1.31 and above. This version supports up to 12 WWNs per HBA and up to 16 WWNs per VC Fibre Channel uplink port. This feature's capabilities will be expanded very soon with an additional firmware update.

#28: Virtualized MAC and WWN (NATing) features on some switches (for example, Cisco's FlexAttach* feature) provide the same benefits as VC's Managed MAC addresses and WWNs

Incorrect: One of the many features provided by Virtual Connect is the ability to “manage” the server blade MAC addresses. Specifically, Virtual Connect ‘manages’ the server blade MAC & WWN addresses. Virtual Connect does not “virtualize” the server blade addresses. Many VC administrators don’t appreciate the difference between ‘virtualized’ addresses and ‘managed’ addresses.

A virtualized MAC address or WWN is an address that is not really owned and used by a physical NIC or HBA. Often, a virtualized address is an address that replaces the real MAC or WWN address of a physical NIC without the server’s knowledge. In other words, the server thinks it is communicating on the network with MAC address X, however, some device (switch) is replacing real MAC address X with a virtual MAC address Y. This process is effectively Network Address Translation (NAT) of the address by the switch or router. Many of the benefits of MAC address management are lost in this type of implementation.

A managed MAC address or WWN, provided by Virtual Connect, is an address that actually is owned and used by a physical server NIC or HBA. Simply put, the server has been assigned, by the administrator, to use a specific address on a specific physical NIC or HBA port. These managed MAC addresses or WWNs appears to the server as the MAC addresses and WWNs that were burned into the physical NICs and physical HBAs at the factory.

The benefits of VC Managed Addresses are:

- **Advanced flexibility and mobility without switch scripting**
Virtual Connect provides consistency and mobility of managed MAC addresses and WWNs within the data center using “server profiles”. A server profile contains the server’s internal identity (server serial number, UUID, BIOS settings, FC boot parameters, etc.) and a server’s external identity (MACs, WWNs, VLAN assignments, and SAN fabric assignments). A server profile can remain assigned to a server blade bay in an enclosure to maintain the internal and external identity of the server constant no matter what hardware is installed in the slot. However, Virtual Connect also allows the movement of the server profile, with the entire server identity, to any server blade bay in any HP blade enclosure across the data center(s). The movement of a server profile is as simple as a few clicks in the Virtual Connect Web UI or the VC CLI. In addition, VCEM can automate the movement of the server profile. In all cases, the movement of a server profile with the server’s identity does NOT require a separate process to apply a reconfiguration script against an external Ethernet or Fibre Channel switch. With Virtual Connect, the movement of a server profile and the server’s identity is transparent to the Ethernet and Fibre Channel switches.
- **WYSIWYG - What You See (on the server) Is What You Get (on the network)**
There is no discrepancy between what the server thinks its MAC address and WWN are and what the external network sees as the server’s MAC address and WWN. Having only one real MAC address or WWN to manage (versus two with virtualized addresses) per port dramatically reduces the complexity of troubleshooting network and SAN related issues.
- **Server application licensing is maintained after hardware changes**
Many server application licensing mechanisms can key off the server’s MAC addresses. If the server’s MAC address changes (replacing a failed NIC, booting server image on a

different physical server, etc.), then the application licensing may require re-licensing using the new MAC address. Virtualized MAC addresses and WWNs do not address this problem. However, VC's use of managed MAC addresses and WWNs does prevent this problem since the server image (OS) will always see the VC managed MAC address and WWN regardless of which physical server the image is running on.

- **No Performance impact on network and storage devices**

Virtualized MAC addresses and WWN can require that a network device (for example, switch) manipulate every frame a server transmits to replace the server's MAC address or WWN with the virtualized address. Also, when the addresses are edited within the frame by the network device, the frames checksum (CRC) has to be recomputed by the network device. The more frames a server transmits, the more work the network device has to do, which can have an impact on the performance of the network or storage switch. Alternatively, VC's use of managed MAC addresses and WWNs means the server actually transmits and receives using the VC managed address. No device on the network (VC or switch) is required to manipulate the server's frames. This results in absolutely no performance impact on the network.

* http://www.cisco.com/en/US/prod/collateral/ps4159/ps6409/ps5989/solution_overview_c22-489466.html

#29: DHCP Option 82 provides the same server redundancy features as Virtual Connect

Incorrect: DHCP Option 82 only ensures that any NIC port connected to a particular switch port will receive a given IP address. This is only good for maintaining a constant IP address for server rip-n-replace. Alternatively, VC's MAC and WWN address management allows a server blade to be replaced, or moved, or added anytime and anywhere within the VC Domain or across multiple blade enclosures. In addition, VC Managed MAC addresses, WWNs, server serial numbers, and UUIDs ensure application licensing, and other OS Image settings, aren't affected by any of these changes.

DHCP Option 82 - simply makes sure that any NIC that plugs into Switch X, Port Y receives IP address Z. This only maintains an IP address for rip-n-replace.

Table 1.	
DHCP Option 82 Does:	<ul style="list-style-type: none"> • Does keep IP the same for any device that connects its NIC to the specific switch port
DHCP Option 82 Doesn't:	<ul style="list-style-type: none"> • Doesn't provide flexibility <ul style="list-style-type: none"> ○ Doesn't allow server administrator alone to move a server to a different slot in the enclosure or across the data center • Doesn't keep the server serial number the same after hardware replacement. <ul style="list-style-type: none"> ○ Can cause problems with application licensing if application uses server serial number as part of license data • Doesn't keep the server UUID the same after hardware replacement. • Doesn't keep the server MAC addresses the same after hardware replacement. <ul style="list-style-type: none"> ○ Can cause problems with application licensing if application uses server's MAC address as part of license data • Doesn't keep the server WWNs the same after hardware replacement. • Doesn't keep Fibre Channel boot parameters the same after hardware replacement • Doesn't keep PXE boot order configuration in BIOS after hardware

	<p>replacement</p> <ul style="list-style-type: none"> • Doesn't allow for pre-provisioning <ul style="list-style-type: none"> ○ Since MAC addresses and WWNs aren't known until the server is received from the server vendor, administrators can't pre-provision the server ○ Without knowing MAC addresses, things like DHCP reservations can't be completed until server is physically in hand ○ Without knowing WWNs, things like WWN zoning, Selective Storage Presentation/LUN mapping & presentation, etc. can't be completed until server is physically in hand
	Reference: See RFC 3046

Virtual Connect – in addition to moving VLAN and SAN assignments with an OS image across any physical blade within the data center, Virtual Connect also manages a server's internal and external identity so that hardware changes are transparent to the OS and to the external LAN and SAN.

Table 2.	
Virtual Connect Does:	<ul style="list-style-type: none"> • Does keep IP the same for any server NIC when using VC MAC addresses and DHCP reservations. • Does provides flexibility <ul style="list-style-type: none"> ○ Does allows server administrator alone to move a server to a different slot in the enclosure or across the data center ○ Keeps the same VLAN assignments with the server ○ Keeps the same SAN assignments with the server • Does keep the server serial number the same after hardware replacement. <ul style="list-style-type: none"> ○ Prevents problems with application licensing if application uses server serial number as part of license data • Does keep the server UUID the same after hardware replacement. • Does keep the server MAC addresses the same after hardware replacement. <ul style="list-style-type: none"> ○ Prevents problems with application licensing if application uses server's MAC address as part of license data • Does keep the server WWNs the same after hardware replacement. • Does keep Fibre Channel boot parameters the same after hardware replacement • Does allow for pre-provisioning <ul style="list-style-type: none"> ○ Since MAC addresses and WWNs are known BEFORE the server is received from HP, administrators can preprovision the server ○ Knowing MAC addresses before receiving the server from HP means things like DHCP reservations can be completed BEFORE server is physically in hand ○ Knowing WWNs before receiving the server from HP means things like WWN zoning, Selective Storage Presentation/LUN mapping & presentation, etc. can be completed BEFORE server is physically in hand
Virtual Connect Doesn't:	<ul style="list-style-type: none"> • N/A – VC provides same result as DHCP Option 82 and provides additional features

#30: VC-FC doesn't provide login distribution and failover on FC uplinks to the SAN

Incorrect: FC login distribution and failover features are provided in VC firmware version 1.31 and above. With these features, server HBA fabric logins can be automatically distributed across all VC-FC uplink ports on the same VC-FC module. Should a port fail or lose link, VC-FC automatically re-logs in the WWN into the fabric on another active VC-FC uplink port from the same VC-FC module.

#31: All VC-FC uplinks have to be connected to the same SAN fabric

Incorrect: Support for multiple fabrics per VC-FC module is provided in VC firmware version 1.31 and above. This feature allows a user to connect each of the four VC-FC uplinks to a different SAN fabric and dynamically assign server HBAs to the desired SAN fabric.

#32: VC implements an immature loop avoidance mechanism

Incorrect: Since the Virtual Connect loop avoidance mechanism is modeled after both NIC Teaming/bonding and after simple layer 2 path redundancy features (for example, Cisco's Flex Link feature*), the mechanism inherits the maturity that these seasoned technologies provide. Add to that HP's years of experience in data center networking and NonStop server architecture that the Virtual Connect developers bring to the product, and the result is a highly reliable HP server identity virtualization and I/O management product that thousands of HP customers have successfully deployed across the globe.

*

http://www.cisco.com/en/US/docs/switches/blades/3120/software/release/12.2_40_ex/configuration/guide/swflink.html

#33: VC Uplink failures require re-convergence on the external network and may cause dropped server sessions

Incorrect: Since Virtual Connect connects to the external network in the same way as a virtual server hypervisor (i.e. STP isn't used to manage Layer 2 redundancy), Spanning Tree re-convergence will not occur as long as the administrator has appropriately configured the external switches directly connected to Virtual Connect. Fundamentally, Virtual Connect failover between VC uplinks behaves the same way as failover between Server NICs in a NIC Team (or NIC Bond) – failover from one uplink (or NIC) to another is transparent to the external network's Spanning Tree and does not require a re-convergence. For a complete discussion of HP's recommendations on configuring switches directly connected to Virtual Connect Uplinks, see the section entitled "Cisco Configuration Guidelines for VC Uplink Ports" on page 21 of the whitepaper "[Virtual Connect for the Cisco Network Administrator](#)".

#34: Cisco's N-Port Virtualization (NPV)* or Brocade Access Gateway** provide all the same advantages as VC-FC

Incorrect: Both Cisco's NPV and Brocade Access Gateway are features that use NPIV (N-Port ID Virtualization) to allow a traditional Fibre Channel switch to operate more as a simple Fibre Channel aggregator instead of as a traditional Fibre Channel switch. This reduces domain ID proliferation. Virtual Connect also supports this feature for reducing domain ID proliferation.

However, neither NPV or Access Gateway provides all the additional features that VC Fibre Channel does. For example, FC boot parameter management in server blade BIOS, server move enablement, support for "Server Profiles" that contain a server's complete identity – MAC & VLAN, WWN & Fabric, serial number, UUID, MACs & WWNs), server adds or replacements without rezoning WWNs or reconfiguring host storage presentation, server pre-provisioning using managed WWNs, and so on.

* http://www.cisco.com/en/US/prod/collateral/ps4159/ps6409/ps5989/solution_overview_c22-489466.html

** http://www.brocade.com/san/pdf/datasheets/AccessGateway_Blade_Server_DS_01.pdf

#35: Cisco VFrame Data Center provides the same capabilities as VC

Incorrect: Virtual Connect and Cisco VFrame Data Center are products that target different problems within a customer's data center. Cisco VFrame Data Center is a network-driven service orchestration solution that enables the coordinated provisioning and reuse of physical and virtualized computing, storage, and network resources from shared pools*. Virtual Connect, on the other hand, is a server identity virtualization and I/O management product that provides an enhanced way to add, remove, upgrade, repair, and move server blades within the data center while minimizing the negative affects on the LANs, SANs, and OS images.

Cisco VFrame Data Center is a higher level management product that interfaces with other data center device management tools and products such as Virtual Connect. In other words, Cisco VFrame Data Center scripting capability could be used to control device management tools such as Virtual Connect via its CLI. Virtual Connect technology would enhance a Cisco VFrame Data Center customer's experience by minimizing LAN and SAN configuration changes that result from Cisco VFrame Data Center's resource orchestration and deployment tasks since server changes within the VC Domain can be transparent to the LAN and SAN.

* <http://www.cisco.com/en/US/products/ps8463/index.html>

#36: VC Ethernet can't be connected to Cisco 6500 switches using Virtual Switching System (VSS)

Incorrect: Virtual Connect Ethernet can connect to a Cisco VSS stack the same as VC Ethernet can connect to any other external Ethernet switch. Both Cisco 6500s in the VSS stack look like a single Cisco 6500 switch – transparent to Virtual Connect. VC Ethernet ports (from the same VC module) in a port channel (LAG) can be connected to two different Cisco 6500s in a VSS stack and the port channel operates as a single port channel (as if the ports are connected between a single VC module and a single Cisco 6500) and traffic is load balanced across all ports. Should either of the Cisco 6500s or 6500 modules fail, the EtherChannel on the VC module will still have connectivity through the alternate Cisco 6500 switch/module.

#37: VC Ethernet does not support Unidirectional Link Detection (UDLD)

Correct: UDLD is necessary to prevent a loop for switch ports using Spanning Tree to manage layer 2 redundancy. VC, like hypervisors such as VMware, doesn't participate in the data center Spanning Tree domain. Therefore, UDLD provides no benefit to VC Ethernet or to hypervisors.

#38: VC Ethernet only provides port counters on uplinks

Incorrect: VC provides port statistics for all Ethernet ports – uplinks, downlinks, and stacking links. The counters VC provides for these ports are the following:

IfInOctets	EtherStatsStatsOversizePkts
IfInUcastPkts	EtherStatsStatsJabbers
IfInNUcastPkts	EtherStatsStatsOctets
IfInDiscards	EtherStatsStatsPkts
IfInErrors	EtherStatsStatsCollisions
IfInUnknownProtos	EtherStatsStatsCRCAlignErrors
IfOutOctets	TXNoErrors
IfOutUcastPkts	RXNoErrors
IfOutNUcastPkts	Dot3StatsAlignmentErrors
IfOutDiscards	Dot3StatsFCSErrors
IfOutErrors	Dot3StatsSingleCollisionFrames
IfOutQLen	Dot3StatsMultipleCollisionFrames
IpInReceives	Dot3StatsSQETTestErrors
IpInHdrErrors	Dot3StatsDeferredTransmissions
IpForwDatagrams	Dot3StatsLateCollisions
IpInDiscards	Dot3StatsExcessiveCollisions
Dot1dBasePortDelayExceededDiscards	Dot3StatsInternalMacTransmitErrors
Dot1dBasePortMtuExceededDiscards	Dot3StatsCarrierSenseErrors
Dot1dTpPortInFrames	Dot3StatsFrameTooLongs
Dot1dTpPortOutFrames	Dot3StatsInternalMacReceiveErrors
Dot1dTpPortInDiscards	Dot3StatsSymbolErrors
EtherStatsStatsDropEvents	Dot3ControlInUnknownOpcodes
EtherStatsStatsMulticastPkts	Dot3InPauseFrames
EtherStatsStatsBroadcastPkts	Dot3OutPauseFrames
EtherStatsStatsUndersizePkts	IfHCInOctets
EtherStatsStatsFragments	IfHCInUcastPkts
EtherStatsStatsPkts64Octets	IfHCInMulticastPkts
EtherStatsStatsPkts65to127Octets	IfHCInBroadcastPkts
EtherStatsStatsPkts128to255Octets	IfHCOctets
EtherStatsStatsPkts256to511Octets	IfHCOUcastPkts
EtherStatsStatsPkts512to1023Octets	IfHCOUmulticastPkts
EtherStatsStatsPkts1024to1518Octets	IfHCOUbroadcastPkts

Unique Features Provided By HP Virtual Connect

Virtual Connect provides many unique features not provided by traditional LAN and SAN switches. These HP designed and engineered features are provided to enhance HP server blade deployment and management for HP customers. These features are not provided by traditional LAN and SAN switches because they don't have server system visibility and configurability like Virtual Connect. The following sections provide a technical discussion of several of these unique Virtual Connect features.

Managed Server Identities

Internal Server Identity

Virtual Connect manages a server's internal identity in order to provide hardware transparency to OS images. If server blade hardware components must be replaced or if an OS image is moved from a server blade to a completely different server blade, the hardware changes will usually negatively affect the OS image (or installed applications)... unless Virtual Connect is used to provide the virtualization layer between the OS and the hardware. Virtual Connect allows server component replacement and OS image movement between any number of physical server blades because Virtual Connect maintains constant the server's internal identity. This internal identity consists of the server's serial number, the UUID, the MAC addresses for all NICs, and the WWNs of all HBAs. Virtual Connect allows the server administrator to define a "server profile" that contains a managed serial number, a managed UUID, managed MAC addresses and managed WWNs. This individual server profile can then be assigned by the server administrator to any physical server blade in the c-Class blade enclosure. If using Virtual Connect Enterprise Manager, the server profile can be assigned to any physical server blade in up to 100 c-Class blade enclosures.

External Server Identity

Virtual Connect manages a server's external identity in order to minimize interruptions for the LAN and SAN administrators. Virtual Connect allows the server administrator to define a unique set of MAC addresses, a unique set of WWNs, selective VLANs, and selective SANs to a specific "server profile". This individual server profile can then be assigned by the server administrator to any physical server blade in the c-Class blade enclosure. If using Virtual Connect Enterprise Manager, the server profile can be assigned to any physical server blade in up to 100 c-Class blade enclosures.

Preprovisioning Using Managed Server Identities

Besides transparency for server blade adds, moves, and changes, Virtual Connect's managed server identities using "Server Profiles" can also significantly reduce the time required for new deployments. Deployment times can be reduced since an administrator can define a server's identity (by creating a Server Profile in Virtual Connect) and can then use the server's identity information to pre-provision the LAN, SAN, OS builds, etc. An administrator could define, for example, 10 server profiles in Virtual Connect for 10 new HP server blades that will be ordered in the future. Each profile will contain the MAC addresses, WWNs, serial number, UUID, etc of each of the 10 HP server blades that will be deployed in the future.

Knowing in advance the reserved MAC addresses and WWNs of all the soon-to-be-ordered HP server blades, the server administrator can go ahead and request DHCP scope reservations using the MAC addresses, request WWN zoning and Selective Storage Presentation (LUN presentation) using the WWNs, pre-build OS images on spare server blade hardware (by assigning the server profile containing the Virtual Connect server serial number, Virtual Connect UUID, etc), etc. In other words, using Virtual Connect's managed server identities allows an administrator to know the internal and external identity of a server before it's even ordered and use that information to pre-provision so that the new hardware is ready to deploy the moment it's received from shipping.

In summary, Virtual Connect manages a server's internal and external identity using a single "server profile" that can allow adding, moving, or replacing HP server blades anywhere within the data center. This managed server identity can also be leveraged for server pre-provisioning to dramatically reduce the time required for new server blade deployments. This complete collection of capabilities is not provided by any traditional Ethernet or Fibre Channel switch nor provided by any other server blade vendor in the industry.

"LAN Safe" Network Connectivity

Virtual Connect Ethernet allows an entire HP c-Class enclosure full of server blades to connect to the external network sharing one or more VC uplink ports and presents itself to the external network in a way similar to how a large virtual server hypervisor (for example, VMware) full of virtual machines sharing one or more NIC uplink ports presents itself to the external network. Since VC presents the entire enclosure to the external network as a "single, large virtual server host", the external network ports need only to be configured the same as they would if connected directly to a physical server hosting multiple virtual machines.

Since VC provides HP server blade network connectivity just like a hypervisor provides virtual server network connectivity, VC doesn't have to support Spanning Tree just like a hypervisor hosts don't have to support it - yet both provide network redundancy and load balancing. Just like a hypervisor hosts, VC provides network redundancy and load balancing features that leverage NIC Teaming/bonding technology instead of configuration-error-prone switch technologies like Spanning Tree. For example, a Spanning Tree configuration error on any single switch in the data center can negatively affect any other connected switch in the network, in addition to all servers connected to the same network. With Virtual Connect, any redundancy and load balancing configuration problems only affect a single blade enclosure – not the entire network.

In addition, Virtual Connect's loop avoidance mechanism allows an entire c-Class enclosure to be connected to the external network without causing loops. Just like a server using multiple NICs with NIC Teaming to connect to the external network doesn't cause a loop, Virtual Connect also won't cause a loop. Aside from using port mirroring or bridging on server NICs, Virtual Connect won't cause a loop even if you purposefully or accidentally mis-configure it. Virtual Connect automatically prevent loops in the case of VC mis-configuration, switch uplinks connected to incorrect VC uplink ports, etc.

Server Adds, Moves, and Replacements are Transparent to LAN & SAN

Unlike DHCP Option 82, Virtual Connect allows not only transparent server replacements (rip-n-replace), but Virtual Connect also allows transparent server additions and transparent server moves across the data center. Since the Virtual Connect Server Profile contains the server's internal and external identity (as discussed above), a server blade can be added to, or removed from, any blade enclosure or moved between blade enclosures without impact to the external LAN and SAN. In other words, because the VC Server Profile maintains a consistent set of MAC addresses and WWNs for the server and because Virtual Connect moves the VLAN and SAN assignments with the Server Profile, any Virtual Connect managed server blade has the flexibility of being added, moved, or replaced anywhere within a VC Domain or across VC Domains without impact on the external LAN or SAN. By no impact, this means that LAN and SAN administrators do not have to change switch port settings (VLAN assignments, WWN zoning, DHCP reservations, etc) whenever server blade changes occur. This frees LAN and SAN administrators to concentrate on design and maintenance of core data center functions.

Summary of the Virtual Connect Capabilities

The following table provides a summary of the Virtual Connect capabilities in comparison to traditional Ethernet and Fibre Channel switches:

Table 3. Summary of Virtual Connect's benefits

	Virtual Connect	Traditional Ethernet or Fibre Channel Switches
Provides LAN & SAN connectivity for multiple server blades	✓	✓
Reduces cables for blade enclosures	✓	✓
Can be configured to allow internal server-to-server communication	✓	✓
Can be configured to segregate server-to-server communication (Private VLANs, separate Layer 2 domains)	✓	✓
Uplinks can be configured as Port Trunks (EtherChannel) & VLAN Trunks	✓	✓
Supports centralized user management	✓	✓
Provides VLAN tagging/trunking on server downlinks and uplinks	✓	✓
Provides redundant and load balanced connectivity for c-Class enclosure to external LAN & SAN	✓	✓
Provides network troubleshooting tools (for example, statistics and port mirroring)	✓	✓
Provides server NIC with "sticky" IP address for simple RIP-n-Replace	✓	✓ (DHCP Option 82)
Provides management GUI and CLI	✓	✓
Provides Layer 3 routing capabilities inside blade enclosure	✗	(varies)
Provides TACACS+/RADIUS support	✗	✓

Provides port level ACLs, and VLAN ACLs	✗	✓
Provides port trunking (EtherChannel) on server downlinks	✗	✓
Provides user configurable QoS features for individual server NICs	✗	✓
No impact on data center Spanning Tree config or VTP config	✓	✗
Allows blade enclosure to represent itself as a single, large hypervisor host to external network	✓	✗
Does not require configuration & management of traditional switches inside each blade enclosure	✓	✗
Maintains consistent UUID for OS image after server hardware changes/moves	✓	✗
Maintains consistent server serial number for OS image after server hardware changes/moves	✓	✗
Maintains consistent MAC addresses after server hardware changes/moves	✓	✗
Maintains consistent WWNs after server hardware changes/moves	✓	✗
When server admin moves server profile, a downlink's VLAN assignments automatically follows	✓	✗
When server admin moves server profile, a downlink's VLAN tagging configuration automatically follows	✓	✗
When server admin moves server profile, an HBA's SAN fabric assignment automatically follows	✓	✗
Maintains consistent PXE setting for server NICs after server hardware changes/moves	✓	✗
Maintains consistent FC Boot parameters after server hardware changes/moves	✓	✗

Enables Server Admin to make server network connectivity changes without interruption of LAN and SAN admins	✓	✗
Device mis-configurations only affect network connectivity for blade enclosure and won't cause problems for external devices on the network (see #3)	✓	✗
Enables LAN, SAN, and OS pre-provisioning before physical servers are ordered/physically received since serial number, UUID, MAC addresses, and WWNs are known ahead of time	✓	✗
Administrator is unable to accidentally or purposefully create a broadcast storm or Layer 2 loop between the enclosure and the external network	✓*	✗
Provides loop-free connectivity out of the box with no user configuration regardless of the external network configuration	✓	✗**
Provides VLAN ID translation (mapping) between tagged server blade NICs and the external tagged network (see #9)	✓	✗
Single management interface (GUI and CLI) for all Ethernet and Fibre Channel modules in blade enclosure	✓	✗

* Server NIC configurations (NIC bridging) and VC port mirroring configurations excluded

** STP or other Layer 2 redundancy mechanisms must usually be specifically configured to guarantee no loops

Additional Resources and References

Virtual Connect Cookbook:

<http://bizsupport.austin.hp.com/bc/docs/support/SupportManual/c01471917/c01471917.pdf>

-or-

www.hp.com/go/bladeconnect (see the Virtual Connect Interest Group)

Virtual Connect Documentation:

www.hp.com/go/bladesystem/documentation

Virtual Connect Firmware:

www.hp.com/go/bladesystemupdates

HP NIC Teaming for Windows Whitepaper:

<ftp://ftp.compaq.com/pub/products/servers/networking/TeamingWP.pdf>

HP Services:

www.hp.com/go/bladesystem/services

BladeSystem Solutions:

www.hp.com/go/bladesystem/solutions

c-Class Port Mappings:

- [c7000 Enclosure](#) (page 8)
- [c3000 Enclosure](#) (page 8 & 9)

About the Author

M. Sean McGee, CCIE #18040, is the senior network architect for the HP BladeSystem Engineering division. Sean began his career at Hewlett-Packard 10 years ago when he joined the Networking Products Division supporting switching and routing products in the 3rd Level support group. Over time, he focused his expertise on data center networking technologies and has spent the last several years in HP engineering groups responsible for HP ProLiant NICs, HP NIC Teaming, HP BladeSystem Ethernet switch architecture, and HP BladeSystem Virtual Connect development. As a member of the HP BladeSystem Engineering division, Sean works with various engineering groups as an internal networking technology consultant for new products, as a technical trainer to customers, partners, and HP Field Engineers, and as a network design consultant to many HP customers deploying HP BladeSystem products.

Appendixes

Appendix A: Frequently Asked Questions

Q1: Why do I see lots of dropped frames (discards) on standby VC uplink ports?

A1: An external switch has no concept of which VC link is the active uplink and which is the standby uplink. As far as the external switch is concerned, one of the uplinks is just a whole lot busier. That means that the external switch is still going to send some types of frames down the standby link and the standby link is going to discard them. This includes all broadcasts, multicasts, and unknown unicasts (destination lookup failure in the CAM table on the external switch). Any of these frames that are received on the standby link will be dropped and will be reflected in the counters.

Q2: Can I manually choose which port channel is the preferred channel for a vNet?

A2: No, VC does not currently support setting the 'port role' for vNets (or Shared Uplink Sets) with LACP enabled (connect mode 'auto'). VC determines which port channel is active and which is standby, based on the following criteria, 1 - Number of functional uplinks per port channel, 2 - Total bandwidth provided by each port channel, 3 - if the previous two are a tie, then the VC module with the lowest MAC address (see TOE tag) will provide the active port channel. In situations where a primary VC port/path fails and is restored, VC automatically fails back only if the restored port/path is better than the current active port/path. This prevents unnecessary failbacks. With the information above, the Administrator could construct the primary port channel with one additional uplink in order to make it the preferred port channel.

Q3: Do I have to use the same load balancing algorithm on both sides of the same port channel?

A3: No, you can have different load balancing algorithms on each side of a port channel.

Q4: I see Link Layer Discovery Protocol (LLDP) frames when I connect a network trace analyzer to a VC uplink port. What is VC doing with LLDP and can I disable it?

A4: LLDP is the IEEE equivalent to Cisco Discovery Protocol (CDP). It is a Layer 2 protocol that allows one device to both announce itself (and some of its feature set) to a neighboring device as well as discover other connected devices on the network. It is extremely low-bandwidth and is unobtrusive. VC uses LLDP to determine when one of its uplinks or cross-connects is directly connected to another VC module in the same VC domain so that it can form a stacking link. There is currently no way to disable it.

Q5: I don't have any CX4 10Gb cables to form stacking links. Can I combine multiple 1Gb RJ-45 links instead?

A5: Yes, by adding multiple 1Gb links between modules, VC automatically aggregates them together to form a single 802.3ad port trunk.

Q6: I am trying to get 802.3ad Port Trunking to work but can't seem to get it to pass traffic. What am I doing wrong?

A6: VC currently only supports LACP for 802.3ad port trunks. Cisco's PAgP is not supported by VC. Future versions of VC may add additional features to enhance this functionality. For LACP to work properly on a Cisco switch, the channel mode must be set to either Active or Passive. VC can display detailed Ethernet and port trunk statistics by clicking the desired Ethernet module under Hardware Overview in the left hand tree-view of VCM.

Q7: How do I setup a cluster heartbeat network in Virtual Connect?

A7: Create a vNet in VC and do not assign a VC uplink port to it. Next, assign a "heartbeat" NIC from each blade in the cluster to this vNet. All heartbeat traffic will be contained within the vNet and will not be transmitted outside of the enclosure.

Q8: I need more than 16 VC uplinks. If I add more VC Ethernet modules to add more uplinks, am I required to use additional NICs on my servers?

A8: No, you can add more VC Ethernet modules, stack them with the other VC Ethernet modules and just use the uplink ports. Any VC uplink on any VC Ethernet module can be used to provide external connectivity for any downlink on any VC Ethernet module.

Q9: I need more NICs on my server blades. If I add more VC Ethernet modules to add more downlink ports, am I required to use additional VC uplinks ports to provide connectivity for these new downlink ports?

A9: No, you can add more VC Ethernet modules, stack them with the other VC Ethernet modules and the new downlink ports can be configured to use the uplinks on the existing VC-Enet modules. Any VC uplink on any VC

Ethernet module can be used to provide external connectivity for any downlink on any VC Ethernet module.

Q10: I noticed that the VC Ethernet module in interconnect bay 1 is the active Virtual Connect Manager and that the VC module in bay 2 is the standby. Does this mean that only the VC module in bay 1 is providing Ethernet connectivity for the server blades?

A10: No. Regardless of which VC module is running the active Virtual Connect Manager, all VC modules can be used simultaneously to provide network connectivity.

Q11: Does VC support iSCSI?

A11: Yes VC is compatible with iSCSI. Since VC is a layer 2 device and iSCSI is an upper layer protocol, above TCP/IP, VC does not implement any features specific to iSCSI. However, VC can provide network connectivity for a server running iSCSI just like any other protocol.

Q12: Why are failovers taking longer than expected to restore connectivity for the server blades (taking 20 seconds or longer)?

A12: Make sure the upstream Cisco switch ports connected to the VC uplink ports are configured with PortFast enabled ("spanning-tree portfast" or "spanning-tree portfast trunk")

Q13: Does VC interact with STP on my network?

A13: No. VC uplink ports look just like server NIC ports (ex. Physical NIC ports on an ESX server) and VC does not support STP on the VC uplink ports.

Q14: Should I expect BPDUs to be sent from the VC uplink ports to my external Cisco switch ports?

A14: No. VC uplink ports do not transmit BPDUs.

Q15: Can I extend any of my L3 routing protocols through the VC domain?

A15: Since VC is a layer 2 device, it does not support any routing protocols. However, layer 3 routing protocols such as OSPF, RIP, RIP2, etc. can be used on the servers and operate transparently through VC.

Q16: Can I configure transmit and receive load balancing NIC teaming with full redundancy throughout the VC domain?

A16: Yes, only if you are using our Integrity blades with INP for Windows/Linux, APA for HP-UX, or Smart Load Balancing in Linux on x86. There are currently no solutions for Windows on x86.

Q17: Can I mix VC 1/10 Gb-F and VC 1/10Gb in the same enclosure?

A17: yep. great way to increase bandwidth while further minimizing cable and port usage

Q18: Does VC support VMware and other OSs that support host-based VLAN tagging?

A18: Yes. See appropriate sections above.

Q19: Can I use third party branded SFPs and XFPs in a VC 1/10 Gb-F module?

A19: No, only HP branded SFP and XFP modules are supported.

Q20: Will upgrading the VC firmware require an outage?

A20: In general, HP recommends upgrading VC firmware during a scheduled maintenance window. However, VC is able to perform a rolling, non-intrusive upgrade of all modules so long as redundancy is configured throughout the solution.

Q21: How fast should I expect a failed VC uplink port to failover to a standby VC uplink port?

A21: 5 seconds or less for an optimized configuration

Q22: How fast should I expect a port channel (LAG) to failover?

A22: 5 seconds or less for an optimized configuration.

Q23: How fast should I expect my teamed NICs to failover when a vNet fails over between uplinks?

A23: 5 seconds or less for an optimized configuration.

Q24: Can I connect VC-Enet's XFP ports to XENPACK or X2 ports on a Cisco switch?

A24: Cisco states that XENPACK, X2, and XFP modules are compatible if using the same port type. Since VC only supports 10GB-LR and 10GB-SR, Virtual Connect should be compatible with XENPACK, X2, and XFP

modules from Cisco as long as they are using 10GBASE-LR or 10GBASE-SR modules. Since Cisco supports several port types, several transceiver types, and many different interface modules for their switches, HP does not test every combination.

http://www.cisco.com/en/US/prod/collateral/modules/ps5455/prod_brochure0900aecd8034bba6.pdf

Q25: Is Virtual Connect compatible with layer 3 protocols other than IP? For instance, does VC support IPv6, IPX, AppleTalk, etc.?

A25: Virtual Connect only supports IP (IPv4) on its management interfaces (Web, SSH CLI, or SNMP). In reference to Virtual Connect's bridging functionality, VC supports any layer 3 or higher protocol in use on server blades. Since Virtual Connect is a layer 2 device, it is layer 3 protocol agnostic. Meaning, the server blades can communicate through VC using any upper layer protocol (for example, IPv4, IPv6, IPX, AppleTalk, etc.) that's carried within an Ethernet frame.

Q26: Does Virtual Connect support jumbo frames?

A26: Yes, VC-Enet supports Ethernet frames sizes up to 9216 bytes.

Q27: Does Virtual Connect support EtherChannel/802.3ad/SLB on the downlinks to the server NICs? Can I use LACP port trunking on the server NICs connected to Virtual Connect?

A27: No, Virtual Connect does not support EtherChannel/802.3ad on the downlinks to server NIC ports.

© 2008 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

VMware and VMware ESX server are trademarks or registered trademarks of VMware, INC. or its subsidiaries in the United States and other countries.

Cisco and EtherChannel are trademarks or registered trademarks of Cisco Systems, Inc. or its subsidiaries in the United States and other countries.

Citrix and Xen are trademarks or registered trademarks of Citrix Systems. or its subsidiaries in the United States and other countries.