

Virtual Connect Multi-Enclosure Stacking Reference Guide

Table of contents

Executive Summary	2
Multi-Enclosure Stacking Guidelines	2
When to choose Multi-Enclosure Stacking	3
Virtual Connect Domain Stack Planning	4
Virtual Connect Fibre Channel	9
Multi-Enclosure Stacking External Connection Architectures	10
Consolidating Existing VC Domains and Enclosures	19
Removing an enclosure from a multi-enclosure domain	20
For more information.....	20



Executive Summary

Customers today are challenged with simplifying the network edge by reducing complexity and cost. HP's Virtual Connect solution helps customers simplify server network connections by cleanly separating the server enclosure from the data center network. Stacking multiple enclosures enables the management of up to four enclosures from a single control point.

ME Stacking provides the following benefits:

- Provide connectivity for any blade server to any uplink port in the VC domain, regardless of location
- Reduce the overall number of cables needed for uplink connectivity
- Ability to move a profile between enclosures
- Reduce the datacenter core switch traffic, as internal communication stays inside the Virtual Connect domain
- Reduces the management touch points by consolidating Virtual Connect manager interfaces

Within environments where the 10Gb uplink ports are a premium, VC Flex-10 and FlexFabric can be used with a minimal number of uplink ports in order to provide connectivity for multiple enclosures. The stack links provide high speed connections between enclosures that can be adjusted by adding more physical links to the stack. This enables the ability to create a high speed private network connection between blades for services such as VMotion, Cluster heart beats, backups, and IP Storage to name a few.

This reference guide will provide best practices to designing a multi-enclosure stack, and provide guidance on various approaches.

Multi-Enclosure Stacking Guidelines

ME Stacking enables a base enclosure to have external physical connections to multiple remote enclosures, by using the physical uplink ports (called Stack Links) of the Virtual Connect (VC) modules. Up to 4 Virtual Connect-enabled C7000 enclosures can be stacked together and managed as one Virtual Connect domain for 128 servers, not to exceed a total of 16 VC-Enet and 16 VC-FC modules (a Virtual Connect FlexFabric module counts as both one VC-Enet module and one VC-FC module towards these limits). Any available uplink port not assigned to a Virtual Connect Network or SUS can become a stack link. All VC Ethernet modules must be stacked together, to form a valid L2 ring topology. Any Ethernet module(s) not part of the entire domain stack is an invalid and unsupported configuration.

The Multi Enclosure Stacking feature was designed and tested for enclosures that are physically close to each other, in a single rack or in a few racks side-by-side. However, stacking can be achieved and is supported with any of the cabling options supported by Virtual Connect, including optical fiber on longer distances. As the distance between enclosures in the stack increases, so will the network latency over the management network (connected through the Onboard Administrators) which will have an impact on Virtual Connect Manager performance, so it is advisable to keep distances between enclosures in a stack to a minimum.

In order to configure ME Stacking, the base enclosure must have 1/10Gb-F Virtual Connect Ethernet Modules or Virtual Connect Flex-10 or FlexFabric Modules in bay 1 or redundant pairs installed in bays 1 and 2. The original HP 1/10Gb Virtual Connect Copper Ethernet Modules¹ are not

¹ HP Part Number 399593-B22

supported in bay 1 and 2 of the base enclosure. The HP 1/10Gb VC Copper module does not contain the required amount of system resources to manage multiple enclosure configurations. Remote enclosures may be populated with any Virtual Connect Ethernet modules. The HP 1/10Gb-F Ethernet Modules can be the management modules of a Multi Enclosure Domain ONLY if the stack does NOT contain any Flex-10 or FlexFabric module, i.e. it only has 1/10Gb and 1/10Gb-F modules (plus VC-FC modules), again because of limited system resources.

Virtual Connect Enterprise Manager is not required to manage a ME Stacked Domain. VCEM provides a central console to administer LAN and SAN address assignments, perform group-based configuration management, and execute rapid deployment, movement and failover of server connections and their workloads for 250 Virtual Connect domains (up to 1,000 enclosures and 16,000 server blades). ME Stacking and VCEM can be used together to provide the stacking benefits listed above and all the benefits of VCEM for central management of all the domains.

The first Virtual Connect domain created on the first enclosure will become the Base Enclosure. Remote Enclosures will then be imported into the Base Enclosure Virtual Connect Domain. The base enclosure will contain the primary and secondary Virtual Connect Manager instances. The base enclosure "role" cannot be transferred to another enclosure in the stack. If the base enclosure (or its VC modules running the management application) is powered off, the other enclosures will remain functional and connected assuming the VC networks have usable uplinks in the remaining enclosures. The Virtual Connect Manager application will not be accessible until the base enclosure is powered back on. All OA's and Virtual Connect modules for each enclosure to be managed must be on the same IP subnet.

It is important that the IP subnet you plan on using for OA and VC traffic is not over utilized. Virtual Connect Manager relies on LAN connectivity between the Onboard Administrators to provide management and control in a multi-enclosure Virtual Connect domain. This connection ensures the configuration settings of the distributed modules remain in sync. HP recommends isolating the management network from other networks to ensure that the communication integrity remains intact.

HP only supports importing remote enclosures that do not have an existing domain configuration. Importing remote enclosures with an existing domain will have its configuration erased. Merging of domains is not an available feature.

Virtual Connect Fibre Channel does not support ME Stacking using the physical uplink port connections. All enclosures in a single stack must have the identical VC-FC and VC FlexFabric module placement and configuration in each enclosure.



Important: Additional guidelines and configuration rules will be found in the document titled "HP Virtual Connect for c-Class BladeSystem Setup and Installation Guide²"

When to choose Multi-Enclosure Stacking

Here are some criteria that can be used to determine if you should use Multi-Enclosure Stacking:

- If the 10Gb ports in the core switches are at a premium, you can reduce the number of ports required to connect your servers by consolidating uplinks over up to 4 enclosures. Of course

² <http://h18004.www1.hp.com/products/blades/components/c-class-tech-installing.html>

the tradeoff is that the overall bandwidth between the enclosures and the core network would be reduced.

- If the applications running on the servers generate a lot of traffic between themselves, but not so much towards the core network, you could use Multi-Enclosure stacking to keep that traffic between up to 4 enclosures without going to the core switches.
- If you have a small environment with 4 enclosures or less and you want to consolidate the management of your Virtual Connect environment without implementing Virtual Connect Enterprise Manager, you can use Multi-Enclosure Stacking to create a single Virtual Connect domain.
- At the opposite end, if you have a large environment with more than 250 Virtual Connect-enabled enclosures, you can use Multi-Enclosure Stacking to group these enclosures in domains of 2, 3 or 4 enclosures to stay below the supported limit of 250 Virtual Connect domains per Virtual Connect Enterprise Manager server.

Virtual Connect Domain Stack Planning

Starting the Multi-Enclosure Stacking setup requires some upfront planning. The base enclosure will need to have its Domain provisioned, before importing any remote enclosures. Imported remote enclosures should not have an existing Virtual Connect domain.



Important: If the remote enclosure contains an existing domain configuration it will be erased.

Follow these steps to configure Multi-Enclosure Stacking through the Virtual Connect Manager GUI.

To import the remote enclosure(s):

- Login to the existing Virtual Connect domain of the first enclosure and select Domain Enclosures in the left pane
- Press the “Find” button
- Enter the IP address and credentials for the OA of the enclosure you wish to import
- Select the enclosure, then press the “Import” button, this may take a few minutes to complete
- From the left pane, select Stacking Links and verify there are no stacking link errors

Figure 1: Importing an enclosure into an existing VC Domain

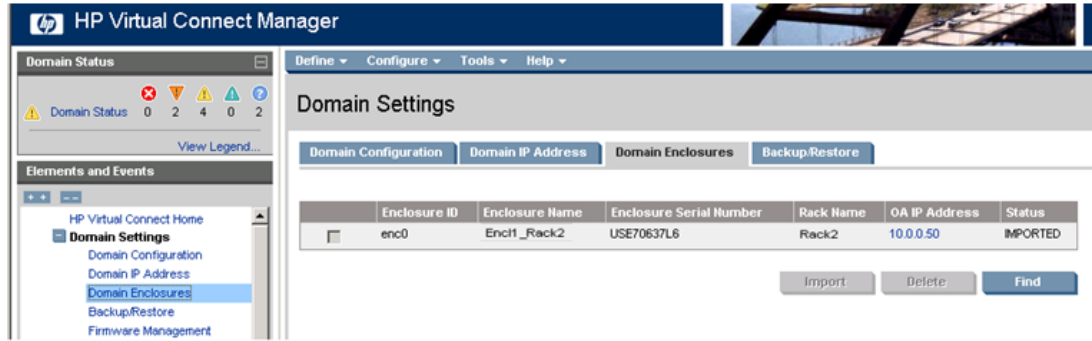


Figure 2: Enter the IP address and credentials of for the OA of the enclosure being imported.

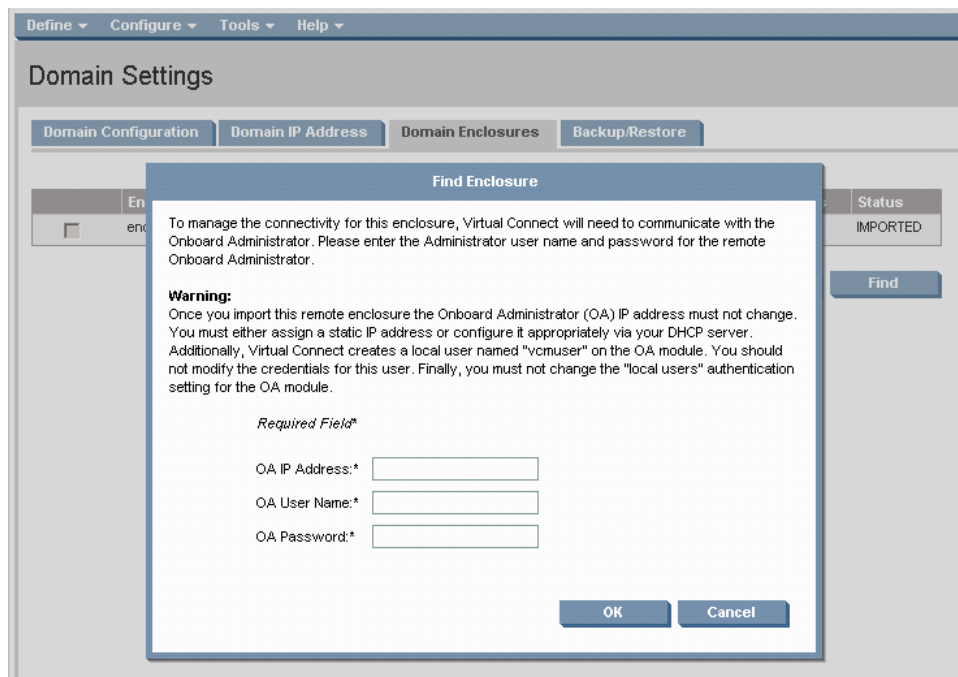


Figure 3: Import the enclosure

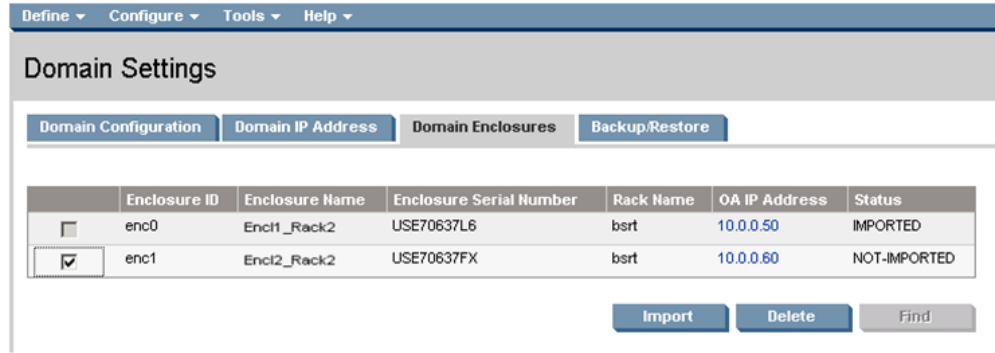
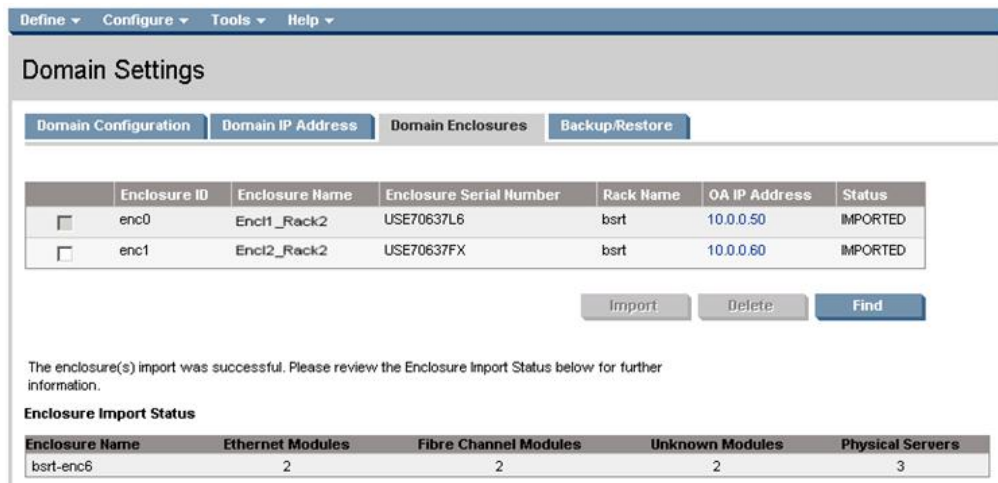


Figure 4: The enclosure is imported



The enclosures in the stack are assigned an enclosure ID. The base enclosure is assigned enc0, and the remote enclosures are given ID's enc1 through enc3. The Enclosure name being imported is also displayed in the second column of the graphic.

Figure 5: VC Domain Stacking Links

Stacking Links

Stacking Links

The Virtual Connect Manager has discovered the VC-Enet module stacking links listed below. The 'Connection Status' below indicates whether all VC-Enet modules are interconnected and accessible. The 'Redundancy Status' indicates whether the VC-Enet modules will remain connected with the loss of a module or cable.

Connection Status: ✔ OK ?

Redundancy Status: ✔ OK ?

Enclosure	Link Speed	From Connection	To Connection
Enc1_Rack2(enc0)			
	10 Gb	enc0:Bay1:PortX1	enc3:Bay1:PortX1
	10 Gb	enc0:Bay1:PortX5	enc1:Bay1:PortX6
	10 Gb	enc0:Bay1:PortX7	enc0:Bay2:PortX7
	10 Gb	enc0:Bay1:PortX8	enc0:Bay2:PortX8
	10 Gb	enc0:Bay2:PortX1	enc3:Bay2:PortX1
	10 Gb	enc0:Bay2:PortX6	enc1:Bay2:PortX6
Enc2_Rack2(enc1)			
	10 Gb	enc1:Bay1:PortX1	enc2:Bay1:PortX1
	10 Gb	enc1:Bay1:PortX7	enc1:Bay2:PortX7
	10 Gb	enc1:Bay1:PortX8	enc1:Bay2:PortX8
	10 Gb	enc1:Bay2:PortX1	enc2:Bay2:PortX1
Enc3_Rack2(enc2)			
	10 Gb	enc2:Bay1:PortX6	enc3:Bay1:PortX6
	10 Gb	enc2:Bay1:PortX7	enc2:Bay2:PortX7
	10 Gb	enc2:Bay1:PortX8	enc2:Bay2:PortX8
	10 Gb	enc2:Bay2:PortX6	enc3:Bay2:PortX6
Enc4_Rack2(enc3)			
	10 Gb	enc3:Bay1:PortX7	enc3:Bay2:PortX7
	10 Gb	enc3:Bay1:PortX8	enc3:Bay2:PortX8

NOTE:
 Port X0 connects to the internal link between horizontally-adjacent VC-Enet modules
 Port X7 and X8 connects to the internal link between horizontally-adjacent Flex-10 enabled VC-Enet modules.

After successful configuration of the ME Stack, you will be able to create a VC network and choose any uplink port in the VC domain. A profile can be assigned to any server blade bay within any enclosure in the Virtual Connect domain.

Figure 6: Virtual Connect network creation

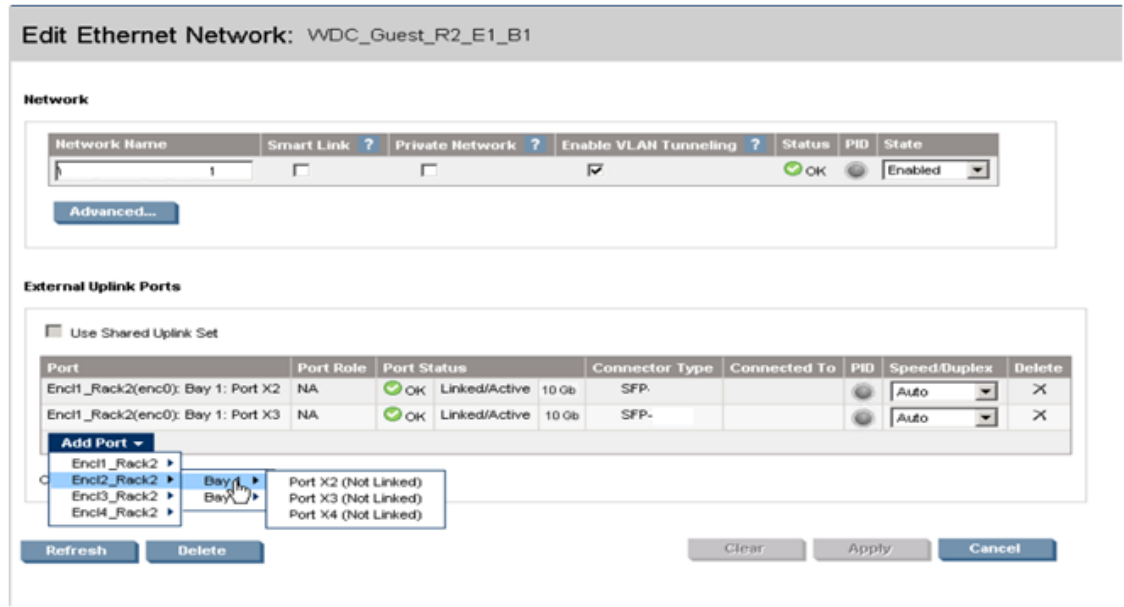
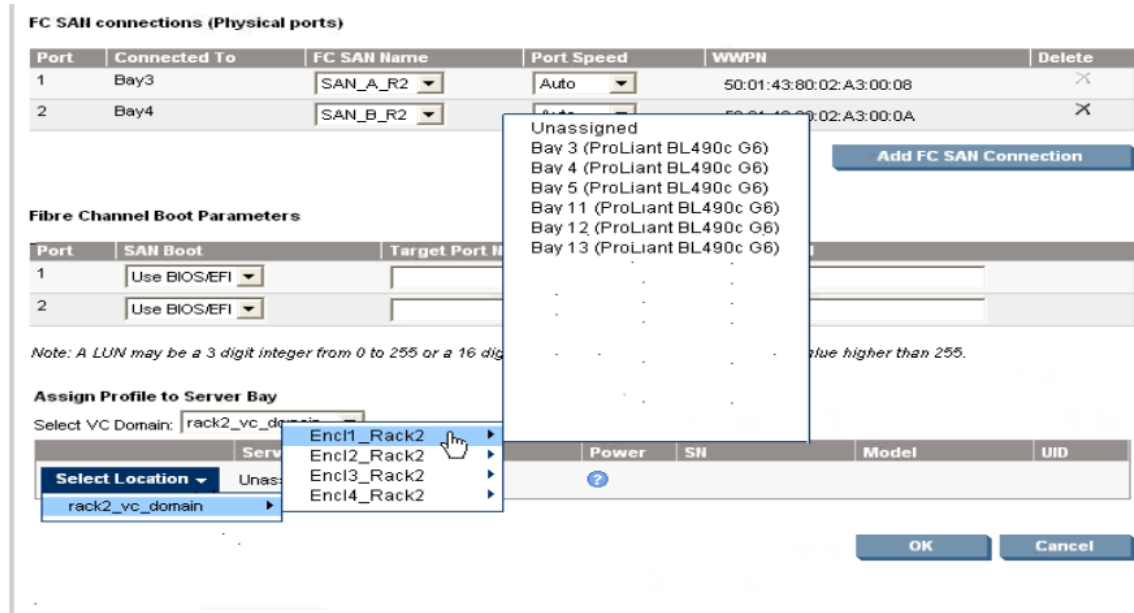


Figure 7: Blade server profile placement



Virtual Connect Fibre Channel

When VC FC or VC FlexFabric is implemented in a multi-enclosure domain, all enclosures must have identical VC-FC and VC FlexFabric module placement and FC cabling. VC-FC does not support the concept of Multi-Enclosure stacking. When a VC Fabric is created, its configuration is automatically replicated to the remote enclosures. For example, if you create a VC FC fabric with a module in Bay 3, port 1 of Enclosure 0 (Encl1_Rack2) which is connected to SAN_A_R2, this configuration will be replicated to the remote enclosures as shown in the graphic below. This helps to ensure that the profile mobility is maintained, so that when a profile is moved from one enclosure to another within the stack, SAN connectivity is preserved.

Figure 8: VC FC SAN Fabric view

Edit SAN Fabric: SAN_A_R2

Fabric

Fabric Name	Status	Login Distribution	Configured Speed
SAN_A_R2	OK	DYNAMIC	Auto

Advanced


Enclosure Uplink Ports

Uplink Port	Enclosure	Bay	Port Status	Speed	Connected To	Delete
Uplink Port 1	Encl1_Rack2	3	OK	8 Gb	10:00:00:05:1E:02:22:2B	X
	Encl3_Rack2	3	OK	8 Gb	10:00:00:05:1E:02:22:2B	
	Encl2_Rack2	3	OK	8 Gb	10:00:00:05:1E:02:22:2B	
	Encl4_Rack2	3	OK	8 Gb	10:00:00:05:1E:02:22:2B	

Add Port

Clear **Apply** **Cancel**

Figure 9: Example of VC FC stacking configuration



Multi-Enclosure Stacking External Connection Architectures

This section will outline multiple scenarios for stacking 2, 3 and 4 enclosures to form a Multi-Enclosure Stack. Each example displayed is configured with 4 uplink ports which connect to the external switches.

The Virtual Connect network would be created with an A and B side to allow all links to be in an active state. Blade host operating systems will need to be configured with NIC teaming to provide both path and module redundancy. More physical uplinks could be setup and additional Virtual Connect networks defined to reduce latency and provide more bandwidth to the networking layer, based on needs and application demands.

The following graphics show the physical connections for stacking enclosures that will provide a configuration that is tolerant of any one module or cable failure. These examples utilize the SFP+ DAC (Direct Attached Copper) connection of each Flex-10 module to interconnect the enclosures.




Important: The CX-4 and X1 ports on the VC Flex-10 modules are shared. If an SFP/SFP+ or DAC cable is inserted into X1, the CX-4 port is disabled. SFP+ DAC cables have a distance limitation of 7m. An alternative to SFP+ DAC cables are the CX-4 cables. They are more rigid and thicker than the SFP+ DAC cables, but offer longer cable distances.



NOTE: CX-4 and SFP+ DAC cables offer a cost effective solution for stacking. Other cabling options supported by Virtual Connect can still be used and are supported.

It is important to note that the following are only recommended examples. It does not pretend to be an exhaustive list of all supported stacking configurations. The basic rules for a supported stacking configuration are listed in the "HP Virtual Connect for c-Class BladeSystem Setup and Installation Guide". The main rule is that all Virtual Connect Ethernet Modules (including FlexFabric Modules) must be interconnected via stacking links. It is also highly recommended to have at least one level of redundancy in the stacking links, meaning that you could lose any one stacking cable and still have connectivity from any VC Ethernet module to any other VC Ethernet module in the domain over the remaining stacking links.

The following examples represent a "dual-ring" topology: the VC Ethernet modules are connected as if they formed 2 rings, one with the left-side modules, one with the right-side modules. Each ring follows the rules laid out above (connectivity from any module to any other even after failure of one stacking cable) independently of the other. This is important in the scenario of a VC firmware upgrade: by default the Virtual Connect Support Utility (VCSU) which is used for firmware upgrades activates the new firmware on all the modules on one side in parallel by rebooting them, then all the modules on the other side after the first side has come back up. Consider a single-ring topology such as this picture with 3 enclosures which has been recommended in the past:



When the modules on the right reboot to activate the new firmware, the servers in the bottom enclosure (enc2) would be isolated from the other 2 enclosures and would lose network connectivity. It is still possible and supported to use such a topology if you are willing to incur this connectivity loss during a firmware upgrade or if you are willing to use the command-line options of VCSU to change the order in which the modules are rebooted to activate the new firmware. Check the VCSU User Guide for the relevant options.

Note that all the modules in a Multi-Enclosure domain must be running the same version of the firmware, it is not supported to run different versions of the firmware in different enclosures that belong to the same domain.



NOTE: When considering an upgrade for Virtual Connect firmware, you should consider upgrades for other components of the BladeSystem c-Class infrastructure. HP recommends the use of Service Packs for ProLiant (SPP) which consist of a comprehensive bundle of firmware and drivers that have been thoroughly tested together. For more information go to <http://www.hp.com/go/spp>


Dual Ring 2 Enclosures

Note that uplinks are connected to alternate enclosures/VC modules (left side/right side). In addition to the external stacking cables (red lines) seen in the following graphic, the VC modules are also cross connected internally, creating a L2 ring topology within Virtual Connect.

This configuration provides the ability to sustain a chassis, module, uplink or upstream switch failure/outage and still maintain connectivity to the network. If additional bandwidth is required, uplinks could be added to these VC networks, or additional VC networks could be configured.

Uplink Ports

Base Enclosure
Enclosure ID enc0



Uplink Ports

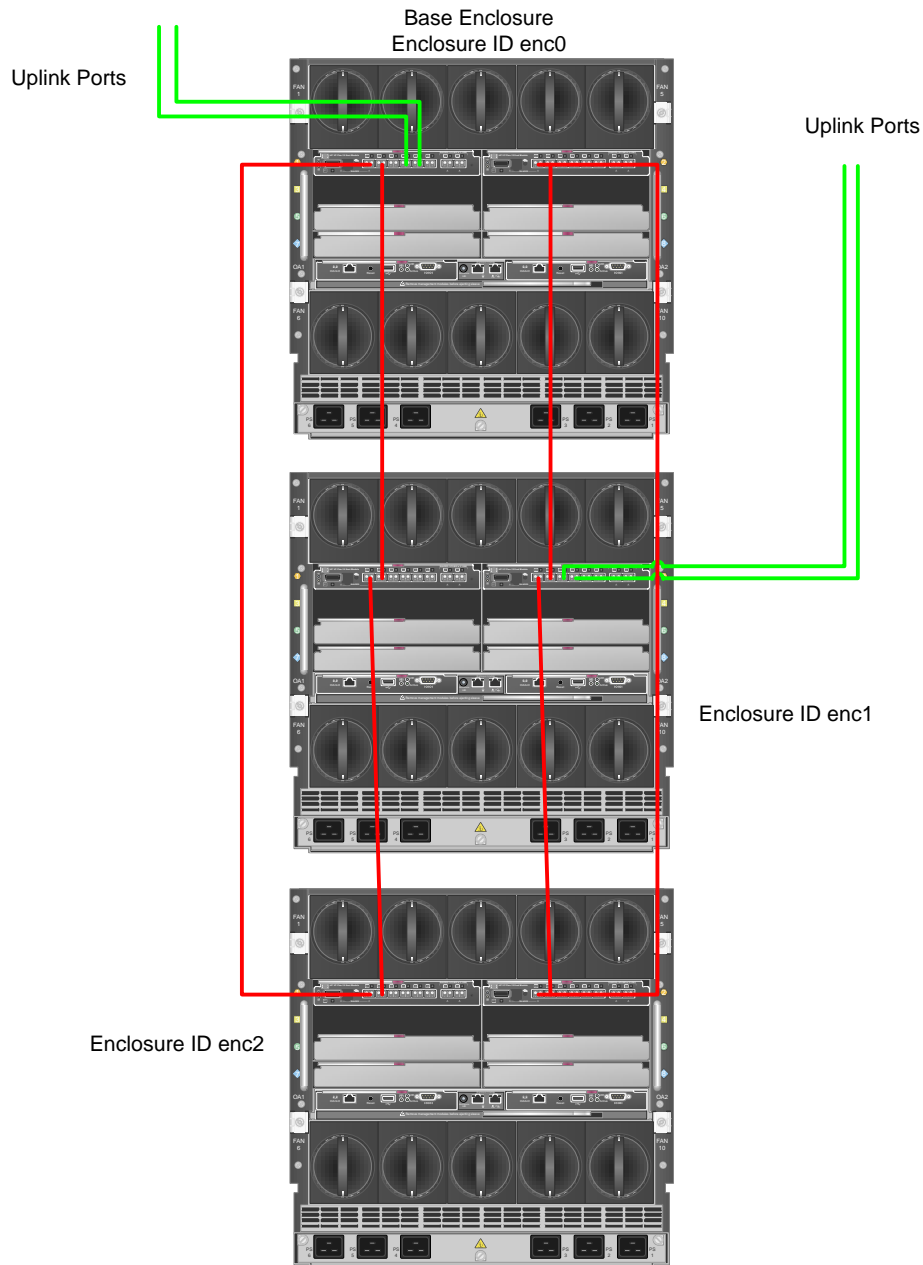
Enclosure ID enc1

Dual Ring 3 Enclosures

Note that uplinks are connected to alternate enclosures/VC modules (left side/right side). In addition to the external stacking cables (red lines) seen in the following graphic, the VC modules are also cross connected internally, creating a L2 ring topology within Virtual Connect.

Communications from Enclosure 3 would be routed through either enclosure 1 or two, depending on how the VC networks are assigned to the server profiles.

This configuration provides the ability to sustain a chassis, module, uplink or upstream switch failure/outage and still maintain connectivity to the network. If additional bandwidth is required, uplinks could be added to these VC networks, or additional VC networks could be configured.




Dual Ring 4 Enclosures

Note that uplinks are connected to alternate enclosures/VC modules (left side/right side). In addition to the external stacking cables (red lines) seen in the following graphic, the VC modules are also cross connected internally, creating a L2 ring topology within Virtual Connect.

Communications from Enclosure 3 and 4 would be routed through either enclosure 1 or two, depending on how the VC networks are assigned to the server profiles

This configuration provides the ability to sustain a chassis, module, uplink or upstream switch failure/outage and still maintain connectivity to the network. If additional bandwidth is required, uplinks could be added to these VC networks, or additional VC networks could be configured.



Additional stacking cables may be required or desired for some of the following reasons:


- If you have more than two VC-Enet modules per enclosure.
- If you want redundancy for more than one cable or module failure.
- If you want to minimize latency between a NIC and its uplink by reducing the hop count.
- If you want to increase bandwidth for IP communication between enclosures

Any of these reasons may lead you to add more physical connections in your stacking environment. You can add additional physical links between VC modules where the first stacking link is located. Virtual Connect will use LACP to form a LAG with each stack link member terminating to the same VC-Enet modules

None of the pictures above show Virtual Connect Fibre Channel modules or Fibre Channel connections since these are irrelevant to the stacking cables layout. Remember that all enclosures in a domain must have the same Virtual Connect Fibre Channel (and FlexFabric) module configuration and the same Fibre Channel uplinks.

Even if the examples above show only 2 Virtual Connect modules per enclosure, the same concepts would apply with more modules. Here is an example with 4 enclosures and 4 modules per enclosure

for a total of 16 modules which would be the maximum supported in a domain:



Choosing where to place uplink cables

There is no specific requirement on where you put uplink cables (the green cables in the example pictures above that connect the VC Domain to upstream switches). It is recommended that you use uplink ports in odd-numbered bays (left side) and even-numbered bays (right side) for redundancy to ensure you retain connectivity during an upgrade where all the modules in the odd-numbered bays reboot to start using a new firmware.

It is customary to put uplink cables on the top-most and bottom-most enclosures as shown in most of the example pictures above. You can also put uplink cables on the first and second enclosures as shown in the 3-enclosure example picture. This has one advantage when you are also running VCEM in the environment: this 3-enclosure domain could coexist in a VCEM domain group with a 2-enclosure domain precisely because the 3rd enclosure does not have any Ethernet uplink, which makes VCEM consider this enclosure as “optional” in the domain group configuration.

You can also put uplinks on more than 2 enclosures, however remember that one of the main benefits of multi-enclosure stacking is to save on the number of uplink cables and ports.

Consolidating Existing VC Domains and Enclosures

This section will outline an example to consolidate multiple independent VC Domains into a ME Stack with one VC domain.

This configuration is more complex as you will be dealing with existing VC domains, VC network uplinks and host configurations. This configuration will need to have some scheduled downtime allotted to make the required changes to the Virtual Connect infrastructure.



Important: ME Stacking will not support “merging” of existing VC domains, remote enclosure VC domains must be deleted in order for stacking to be successful.

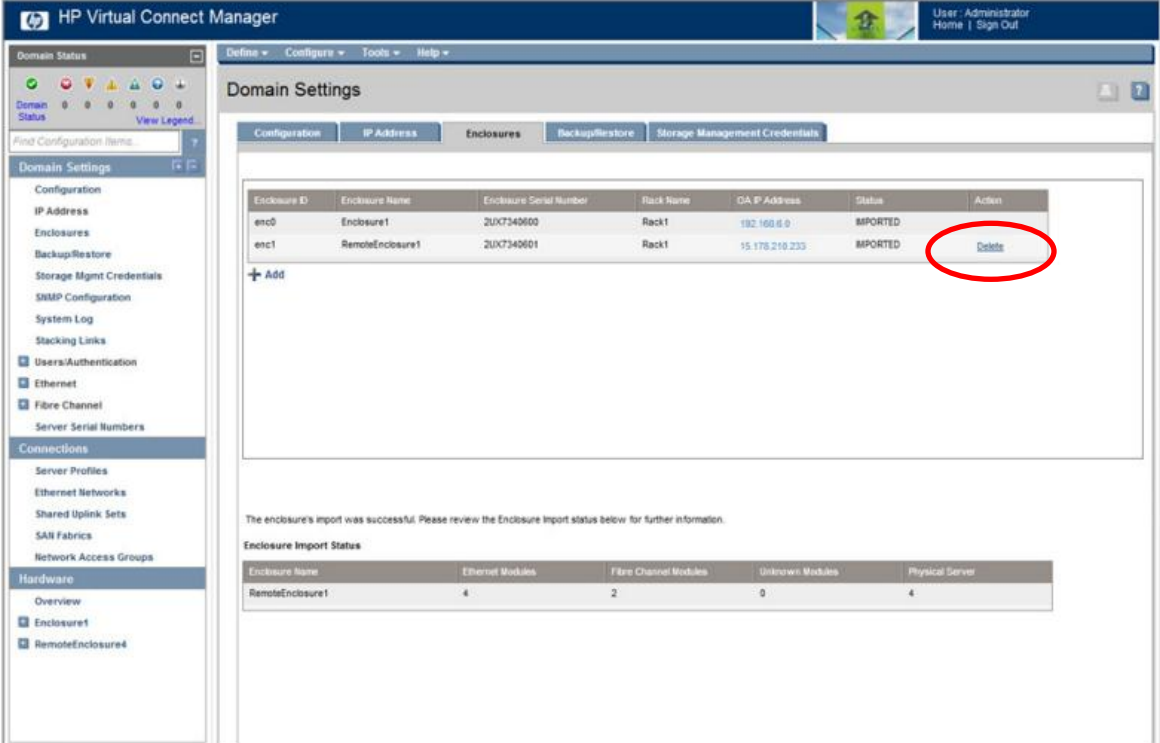
Enclosures being imported must have no previous Virtual Connect domains configured. The Base enclosure becomes the only VC domain for the entire ME Stack.

This is an example set of steps for ‘importing’ a previously configured remote Virtual Connect domain and what you have to plan for.

- Backup ALL VC domains
- Create a VC CLI script to recreate the Profiles for all the blades in the remote enclosures: do not assign them to server bays at this time. The new profile MAC and WWPN’s will be created with the HP Predefined managed range of the base enclosure. If you are using Virtual Connect firmware 3.30 or later, you can use the CLI command “show config” to create a script that can be edited and used to recreate these profiles.
- Decide on new stack uplink port strategy – will all the existing uplink ports be used in the new stacked domain or will additional be added?
- Decide on new naming convention for remote enclosure VC network uplink ports (if the same name as the base enclosure)
- Import the remote enclosures through the CLI or GUI interface.
- Take the new WWPN’s created in the new unassigned profiles and have them zoned and presented in parallel with the existing WWPN’s for storage. Once the enclosures are stacked and the new profiles assigned to the remote servers, remove the old WWPN’s from the storage infrastructure.
- If Needed – Create a CLI script to define new VC Networks using the remote enclosure uplink ports.
- Create a CLI script to assign the new profiles to the correct blade server bays.

Removing an enclosure from a multi-enclosure domain

You can remove an enclosure from a multi-enclosure domain as long as it is not the base enclosure with the VC modules that manage the domain. In order to do so, you must first unassign or move all the server profiles that were assigned to server bays in the enclosure you want to remove. You must also make sure that no network, shared uplink set or port monitor use any uplink port on the VC modules in that enclosure. You can then go to Domain Settings in Virtual Connect Manager and click Delete in the Action column of the row displaying the enclosure you want to remove from the domain. You can also select the enclosure you want to remove, right-click and choose delete from the pop-up menu. After the enclosure is removed, its Virtual Connect modules will be back to unconfigured mode and you will have to create another VC domain in that enclosure (or import it into another multi-enclosure domain).



The screenshot shows the HP Virtual Connect Manager interface. The main window is titled "Domain Settings" and has tabs for Configuration, IP Address, Enclosures, Backup/Restore, and Storage Management Credentials. The "Enclosures" tab is active, displaying a table with the following data:

Enclosure ID	Enclosure Name	Enclosure Serial Number	Rack Name	OA IP Address	Status	Action
enc0	Enclosure1	2UX7340600	Rack1	192.168.6.0	IMPORTED	
enc1	RemoteEnclosure1	2UX7340601	Rack1	15.178.210.233	IMPORTED	Delete

The "Delete" button in the Action column for the "RemoteEnclosure1" row is circled in red. Below the table, there is a message: "The enclosure's import was successful. Please review the Enclosure Import status below for further information." This is followed by an "Enclosure Import Status" table:

Enclosure Name	Ethernet Modules	Fibre Channel Modules	Unknown Modules	Physical Server
RemoteEnclosure1	4	2	0	4


For more information

www.hp.com/go/ bladesystems

Share with colleagues     

Become a fan on  >>

Follow on  >>



Get connected
www.hp.com/go/getconnected
Current HP driver, support, and security alerts
delivered directly to your desktop

© Copyright 2012 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.

4AA1-1555ENW, Created April 2010, Updated April 2012 (Version 3.0)

