



Discover2013

It's time to build a better enterprise.
Together.

TB2603 - Building VMware vSphere 5.1 with blades, Virtual Connect and EVA

Yury Magalif, MASE, VCP

Principal Architect – Cloud Computing

okzebra@gmail.com

yury.magalif@cdillc.com

Thank you!

- » Fifth Year
- » Survey Ratings appreciated!
- » Criticisms
- » Too long
 - » Reduced content
- » No proof numbers/charts
 - » Only 60 min – focus on practical
 - » Take with a grain of salt – Use Scientific Method

Twitter Experiment:

- » Please live Tweet points you find interesting, using the following hashtag:

#HPtrick

- » Look for suggested tricks in the slides.
- » Use this hashtag to chat with me on Twitter:

June 18, 2013 -- Tuesday, 2-3 pm EST

Agenda

- I. Design decisions
- II. Firmware Updates
- III. Virtual Connect
- IV. EVA (P6000)
- V. VMware ESX/ESXi 5.1 (vSphere)
- VI. Summary
- VII. Question & Answer

I. Design Decisions

- » Goal: Virtualize your Infrastructure with VMware
- » Why VMware?
- » Industry leader – 54% of the market per IDC

1. Why Blades & Which blades?

- » Why NOT Blades – no longer a question -- Michael Jordan
- » Run HP Proliant Sizer for VMware
- » Run VMware Capacity Planner analysis
- » Considerations
 - » VMware servers used to be beefy, for extra I/O options.
 - » Now -- Flex-10, can have many NICS in limited physical space
 - » Sweet spot at 192 GB of RAM per blade
 - » WS460c Gen8 Workstation Blade – 8 GPUs, local SSDs, supports VMware Horizon View 5.2 for Virtual Desktop GPU offload – can do AutoCAD
 - » **#HPtrick** Use hot-swap SSD local drives for VMware Host Cache - even for boot-from-SAN blades

2. Which storage and how much?

- » Choose on BOTH capacity (TB) and performance (IOPS)
- » HP All-in-One and Storage Servers (NAS & iSCSI)
 - » For really, really small business
- » HP P2000/MSA array (FC, SAS or iSCSI)
 - » Departmental, decent small business, some midsize business
 - » You will outgrow the P2000 in a couple of years
 - » StoreVirtual 4000 has more features



Which storage and how much, continued

- » HP StoreVirtual 4000 (iSCSI and FC)
 - » Mid-size to large enterprises. A very popular iSCSI/FC solution with unique software features (HA and DR in the box, consolidate local storage with VSA, VMware VAAI integration - plugin).
- » HP P6000/EVA array (FC & iSCSI)
 - » Mid-size and large enterprises.
 - » Long in place upgrade path from entry level
 - » iSCSI, Thin Prov, Dynamic LUN migration, VAAI – in box
 - » Replication and snapshots
- » 3PAR StoreServ – cheaper than P9000/XP, can direct connect to VC FlexFabric, can upgrade from EVA using “Peer Motion” technology.
- » HP P9000/XP array (FC & iSCSI)
 - » Super reliable, redundant, but more expensive than 3PAR.

3. iSCSI or Fibre Channel (FC)?



- » iSCSI taking over the market
- » Cheaper than FC, but NOT cheap
- » Easier to setup for SAN newbies
- » Can use some existing infrastructure
- » Fast, with 10Gbit Ethernet
- » Start with iSCSI, unless brokerage house or already have FC

4. Boot from SAN or local?

- » Advantages
 - » Easiest rip-and-replace for broken blades
 - » Can take advantage of cloning, snapshots and replication of the OS drive on the array – use as a sort of imaging tool.
 - » Host Cache
- » Disadvantages
 - » Server provisioning complex
- » **Boot from SAN recommended, with SSD local drives for cache**



5.1 Blade Interconnect – Virtual Connect, Procurve, Cisco OR Brocade?

- » Cisco & Brocade’s “dumb down” gateway modes are inferior to VC FC WWN replacement and Profiles
- » Cisco’s 3120G – 10Gbit Uplinks, Layer 3; Procurve 6120 – Layer 2, 10 Gbit Up/Downlinks; HP Gbe2c – Layer 2/3 1Gbit. All have many more Ethernet features than VC Ethernet modules
- » VC Ethernet Flex-10 module can program the Flex-10 & FlexFabric adapters to present multiple Physical Functions to the server as standalone PCI devices – like virtual NICs/HBAs.

5.2 Blade Interconnect – Virtual Connect, Procurve, Cisco OR Brocade?

- » FlexFabric 10Gb/24 module is unique with FCoE, dual personality FC/Ethernet ports
- » New Flex-10/10D – if you do 10Gbit iSCSI to your storage & 10Gbit dual-hop FCoE & 10Gbit to your network
- » **#Hptrick** To get multi-hop FCoE with Cisco Nexus 5K switch & HP Blades, use Cisco Fabric Extender for HP, model B22HP.

6.1 Recommendation -- Virtual Connect or Switches?

- » Who do you want to manage VC FC and Eth modules? If Server Admin, get VC
- » If you got VC Eth, get VC FC. Do not mix VC with switches.
- » Do you need MAC address and WWN replacement? You need VC.

6.2 Recommendation -- Virtual Connect or Switches?

- » Do you need to have many virtual, speed flexible NICs? Get VC Flex-10
- » Go with FlexFabric 10Gb/24 port or Flex-10/10D , unless price is a problem – they are the future
- » Do you need Layer 3 routing, VSANs, centralized switch management? Get FC/Eth switches.

7. VMware standard vSwitches, Distributed vSwitch, Nexus 1000v, HP FlexFabric Virtual Switch 5900v?

- » What's the licensing cost? Standard Vswitch is least expensive, then DVS, then Nexus (most expensive for the Advanced edition, but Essential Edition is free)
- » Do I have HP 5900AF Top-of-rack (ToR) access layer physical switch? Get HP FlexFabric Virtual Switch 5900v (Q4 2013 availability)
- » Do I have a Cisco department that refuses to let VM admins manage the network? Get Nexus for Data, DVS for management.
- » Do I want more redundancy and less configuration for my ESX networking? Get DVS.

II. Firmware updates

» Goal: Fulfill
Virtual
Connect
firmware
prerequisites

Warnings on Firmware

- » DCC – changing FlexNIC/HBA port settings without power on/off and SmartLink did not work with ESX 4.1
 - » VC 2.30 (or later) and NC532m or NC532i adapter firmware 2.2.6. w/ ESX 4.1 & Broadcom bnx2x VMware ESX Driver 1.54
- » Any VLAN Tags would disable network functionality with ESX 5.0
 - » Emulex Firmware earlier than 4.0.360.15 with ESXi 5.0 Emulex be2net driver 4.0.88.0 (comes in box with ESXi5)
- » External FC Switches must be at
 - » Brocade FOS 6.1.1b, Cisco SAN-OS 3.3(2), NX-OS 4.1(1c)
- » External Ethernet switches – all current firmware is compatible

Check firmware dependencies in HP SPOCK streams

Virtualization

- » SVSP
- » MS Virtual Server 2005
- » HP Virtual Machines

Other Hardware

- » 3PAR
- » Converged Application Solutions
- » Disk Encryption
- » External Storage
- » Host Bus Adapters
- » iSCSI / FCoE / FCIP / DM (HP)
- » JBOD
- » Mainframe Connectivity
- » Nearline Storage
- » Storage Servers (NAS)
- » StoreVirtual / LeftHand
- » **Switches**
- » Virtual Connect

- » Need HP passport
- » <http://h20272.www2.hp.com/>
- » Look for Switches and VC link on bottom left
- » Get these files
 - » B-Series/C-Series FC Switch Connectivity Stream.pdf
 - » B-Series/C-Series FCoE Switch Connectivity stream.pdf
 - » HP Virtual Connect FlexFabric 10Gb/24port Module for c-Class Blade System
 - » HP Virtual Connect 8Gb 24-Port Fibre Channel Module for c-Class Blade System
 - » HP Virtual Connect 4Gb / 8Gb 20-Port Fibre Channel Module for c-Class Blade System
- » Look for VC sections

HP VC Flex Fabric 10Gb/24-Port Module for c-Class Blade System (571956-B21)

- **VC-FF firmware 4.01**, 3.75 is supported with 5.2(6a), 5.2(6b), and 5.2(8) in an HP SAN.
- VC-FF firmware 3.51, 3.60, 3.70 supported with NX-OS 5.2(1), 5.2(2), 5.2(2d), 5.2(6a), 5.2(6b), and 5.2(8) in an HP SAN.
- VC-FF firmware 3.30 supported with NX-OS 5.0(4b), 5.0(4d), 5.2(2), and 5.2(2d) in an HP SAN.
- VC-FF firmware 3.18 and 3.17 supported with NX-OS 5.0(4b) in an HP SAN.
- VC-FF firmware 3.15 supported with NX-OS 5.0(4b) in an HP SAN.
- The following C-series FC switches are supported: MDS 9222i, SN8000C, MDS 9513, MDS 9509, MDS 9506, MDS 9124, MDS 9134, SN6000C - MDS 9148.

- VC-FC2 firmware 1.04, 1.03, 1.02 & 1.01 supported with NX-OS 5.0(4b), 5.0(4d), 5.2(1), 5.2(2), 5.2(2d), 5.2(6a), 5.2(6b), and 5.2(8)

Check release notes for latest VC firmware

The screenshot shows the HP Support & Drivers website. At the top, there are three tabs: 'Drivers & Software', 'Product Support & Troubleshooting', and 'Support Community'. Below the tabs is a search bar with the text 'virtual connect' and a 'SEARCH' button. Below the search bar, there is a message: '12 matches found. Please select one below'. Below this message, there is a list of search results, with the first one highlighted: 'HP Virtual Connect 4Gb Fibre Channel Module for c-Class BladeSystem'.

» <http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&lang=en&cc=us&docIndexId=64180&taskId=101&prodTypeId=3709945&prodSeriesId=4144084>

» **#HPtrick** If there is a conflict between HP SPOCK streams and Release Notes, follow HP SPOCK

* RECOMMENDED * HP BladeSystem c-Class Virtual Connect Firmware, Ethernet plus 4/8Gb 20-port and 8Gb 24-port FC Edition Component for Windows	4.01 10 Jun 2013
* RECOMMENDED * HP BladeSystem c-Class Virtual Connect Firmware, Ethernet plus 4/8Gb 20-port and 8Gb 24-port FC Edition	4.01 10 Jun 2013
* RECOMMENDED * HP BladeSystem c-Class Virtual Connect Firmware, Ethernet plus 4/8Gb 20-port and 8Gb 24-port FC Edition	3.70 4 Sep 2012

Order of Upgrade

1. External Switches
2. Onboard Administrator
3. iLOs on blades
4. Fibre card
5. NIC
6. Virtual Connect modules

How to Upgrade

- » Service Pack for ProLiant 2013.02.0 with HP Smart Update Manager (SUM) 5.3.5 inside – firmware and drivers all-in-one:
 - » http://h18004.www1.hp.com/products/servers/service_packs/en/index.html
- » Virtual Connect Support Utility (VCSU) v1.8.1 - Available at hp.com/bizsupport, search
- » Use “HP BladeSystem ProLiant Firmware Management Best Practices Implementer Guide”, get it here
 - » http://h20000.www2.hp.com/bc/docs/support/SupportManual/c02049593/c02049593.pdf?jumpid=reg_R1002_USEN

```
C:\Program Files\Hewlett-Packard Company\Virtual Connect Support Utility\vcshp  
HP BladeSystem c-Class Virtual Connect Support Utility  
Version 1.7.1 (Build 21)  
Build Date: Nov 23 2012 05:51:02  
Copyright (C) 2006-2012 Hewlett-Packard Development Company, L.P.  
All Rights Reserved  
Please enter action ("help" for list): healthcheck  
Please enter Onboard Administrator IP Address: 10.0.0.21
```



Firmware 4.01 – New Features


- » **Dual-Hop FCoE Support -- Extending convergence for Flex-10 and FlexFabric**
- » **Priority Queue QoS -- Prioritizing critical application traffic**
- » **Bandwidth Optimization -- Min/Max for NIC Bandwidth**
- » **SNMP MIB enhancements -- Improved monitoring and troubleshooting**
- » **Multicast Filtering -- Single Source Multicast**
- » **Custom role-base configuration -- Delegate management permissions**
- » **GUI & CLI enhancements -- Session timeout & CLI tab completion**

Dual-Hop FCoE

Today

One-Hop

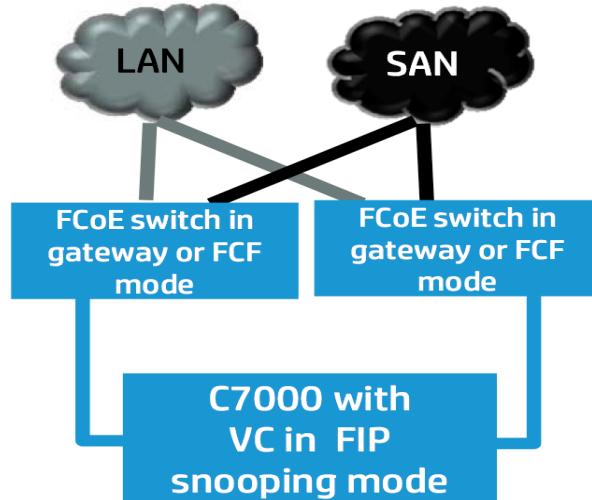
Split FC & Enet
at Enclosure edge



4.01

Dual-Hop


Split FC & Enet
at ToR FCoE switch



post 4.01


Multi-Hop

Connect to upstream
FCF switches

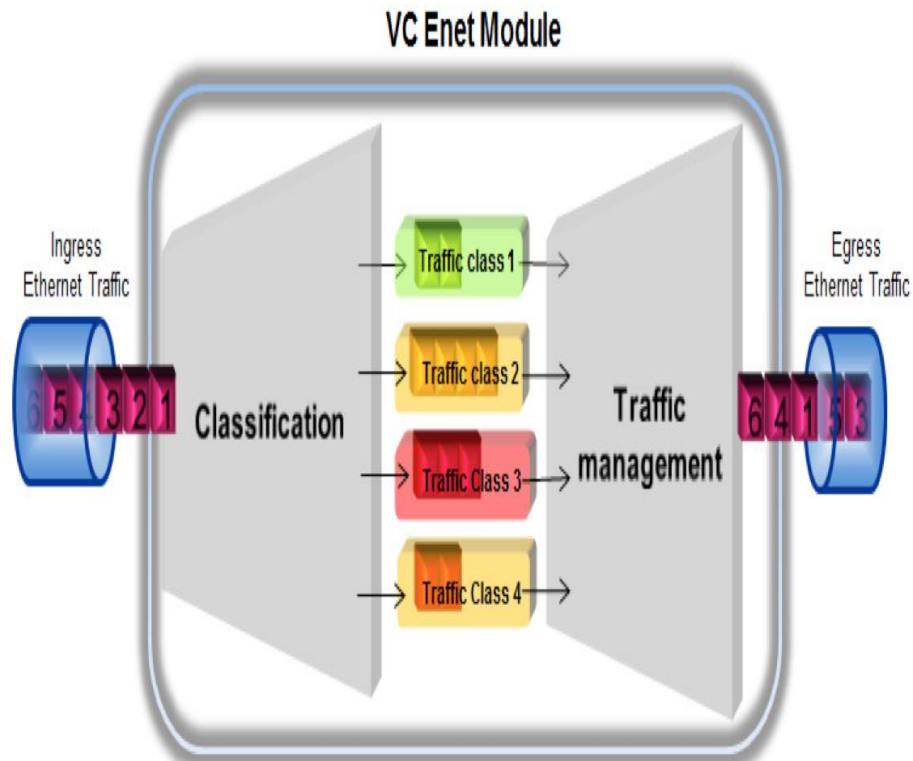


Priority Queue QoS

Traffic pre-4.01 or QoS pass-through



Traffic with QoS



III. Virtual Connect Design & Configuration for VMware

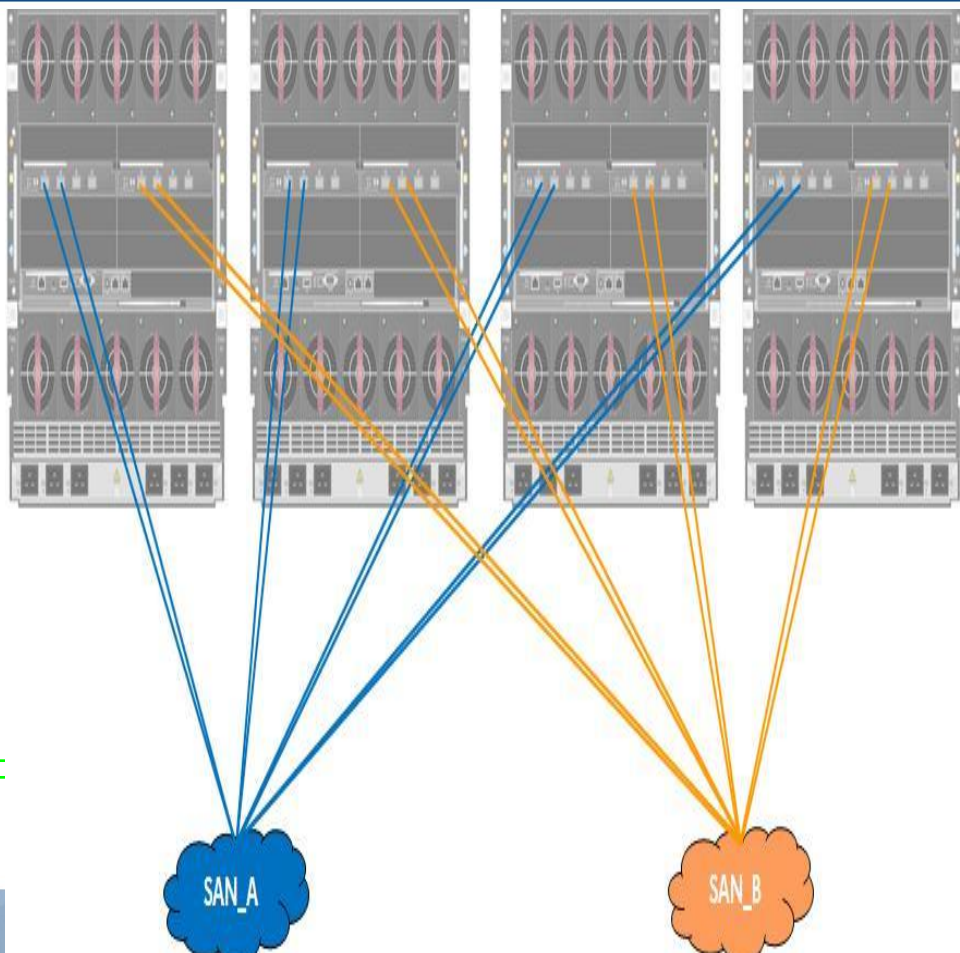
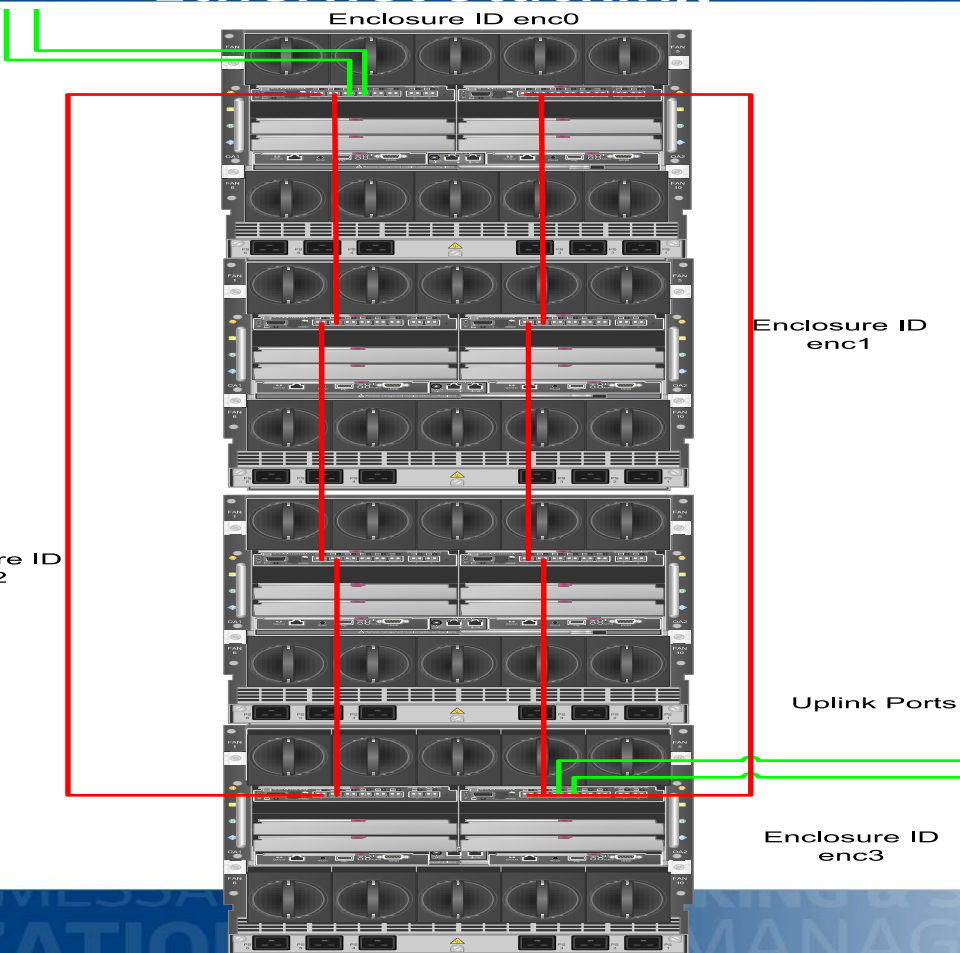
» Goals: Minimize management, get best HA, redundancy & load balancing

1. General design tips

- » If you have multiple enclosures with VC modules, import the 2nd, 3rd & 4th enclosure into the 1st enclosure's VC domain
- » Connect two CX4/SFP+ DAC stacking cables between VC Eth modules as if they formed 2 rings, one with the left-side modules, one with the right-side modules.
 - » Cannot do FC stacking
- » Use HP Default generated values for VC Assigned MAC addresses, WWNs & Serial #s.
 - » If you have multiple enclosures with 1 domain, then use the 1st range (best). If you have multiple enclosures with 1 domain per enclosure, use a different range for each enclosure (HP defined 2, 3, 4 etc.)
- » Name VCnets, Shares Uplink Sets, VC fabrics, Profiles by:
 - » Type, Enclosure, Bay, Where Connecting, Blade model
 - » Ex. vcnet_enc01bay02_pch02, vcfab_enc01bay03_vsan3, vcnet_enc01bay02_vlan20, vcprof_enc01bay01_bT480c

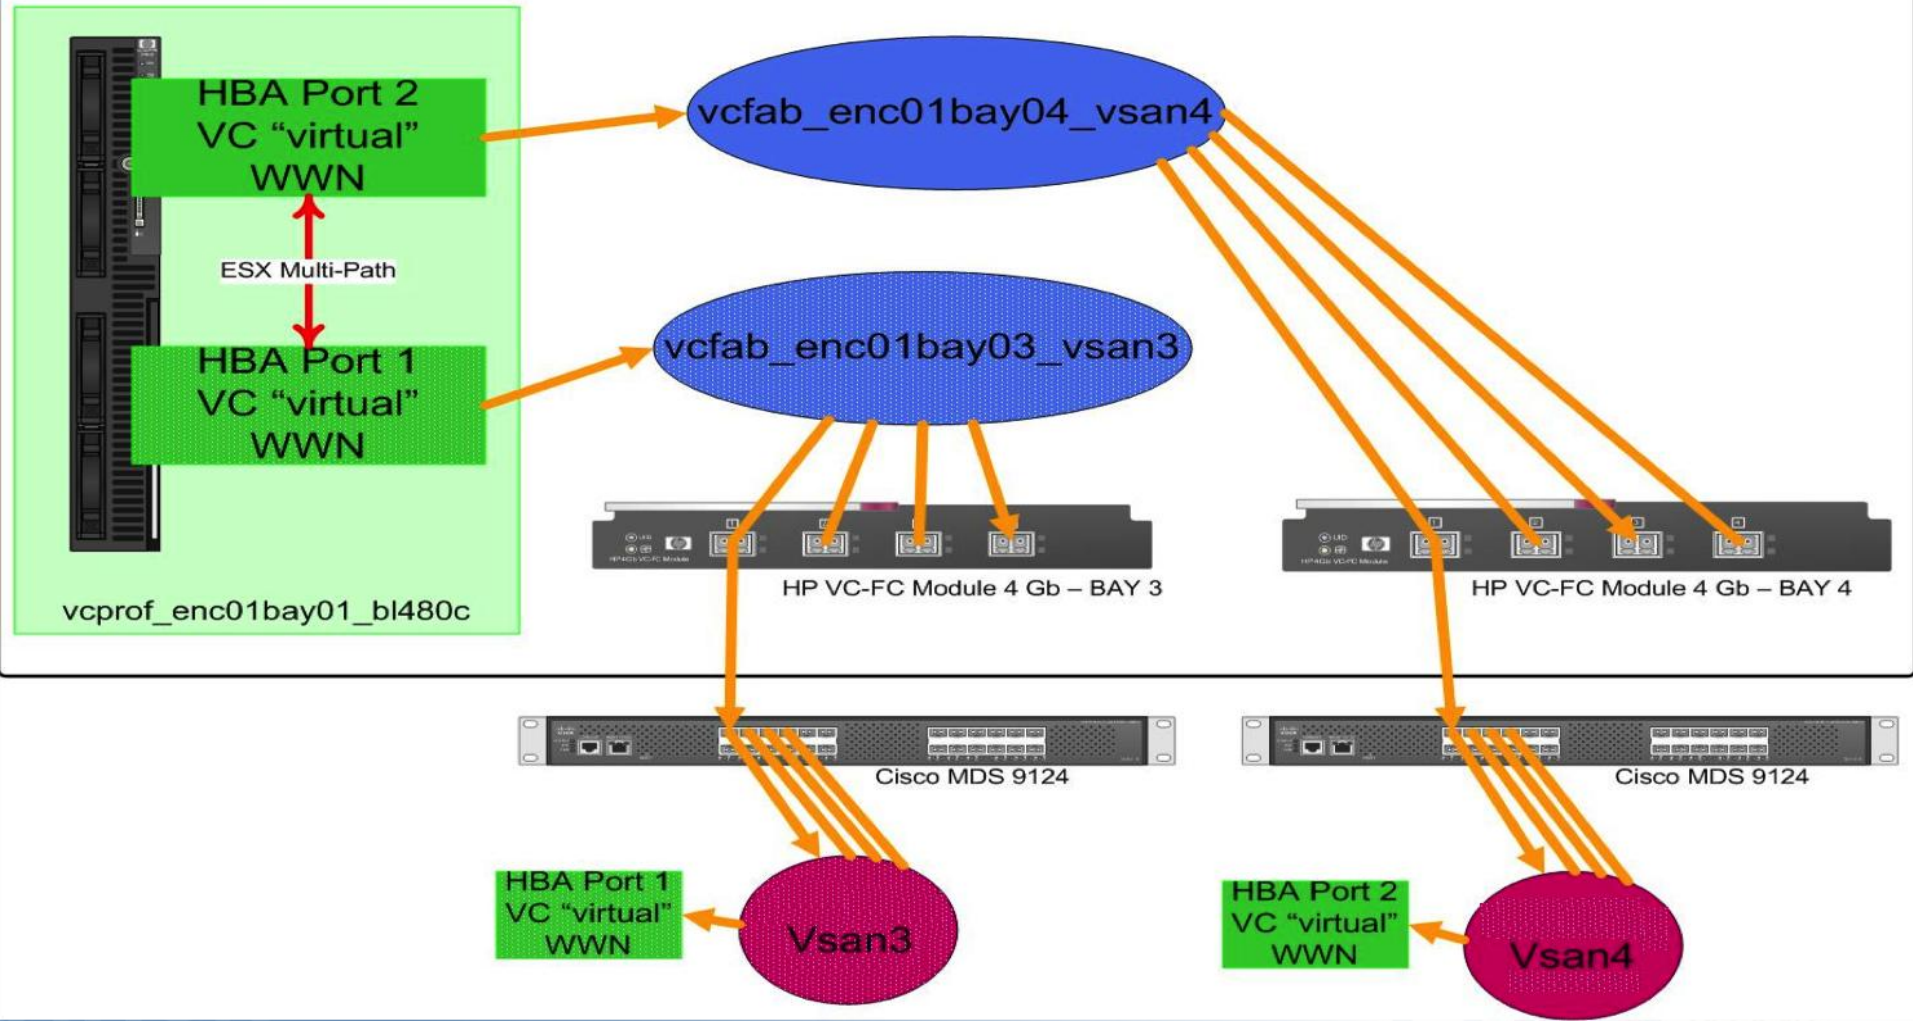
4 encl. 1 VC Domain Ethernet Stacking

No FC Stacking



2. FC Design

- » 1 VC Fabric per VC FC module/bay with 4 external connections grouped
 - » For 8 Gbit 24-port, can group 8 connections, or leave some for dedicated backup network
- » Each module/bay's VC fabric is plugged into a separate SAN switch, for 2 redundant SAN fabrics
- » Add bandwidth by plugging in cables.
- » For FCoE, can do up to 6 FlexNics and 2 FlexHBAs for integrated card (1 per chip), more for added mezzanine cards
- » Can now do northbound FCoE uplinks OR separate FC connections and Ethernet



#HPtrick FlexFabric – must create SAN Fabric 1st, assign the uplink ports, then choose either FCoE or iSCSI – NOT both

Find Configuration Items... ?

Domain Settings

- Configuration
- IP Address
- Enclosures
- Backup/Restore
- Storage Mgmt Credentials
- SNMP Configuration
- System Log
- Stacking Links
- Users/Authentication
- Ethernet
- Fibre Channel
- Server Serial Numbers

Connections

- Server Profiles
- Ethernet Networks
- Shared Uplink Sets
- SAN Fabrics
- Network Access Groups

Hardware

- Overview
- Enclosure 1
 - Interconnect Bays
 - Device Bays

Profile

Profile Name: My_FCoE_Profile Network Access Group: Default Advanced Profile Settings

Ethernet Adapter Connections

Port	Network Name	Status	Port Speed	PXE	MAC
1	net1	✓	PREFERRED	USE-BIOS	VC-DEFINED
2	net10	✓	PREFERRED	USE-BIOS	VC-DEFINED

+ Add

ISCSI HBA Connections

FC HBA Connections

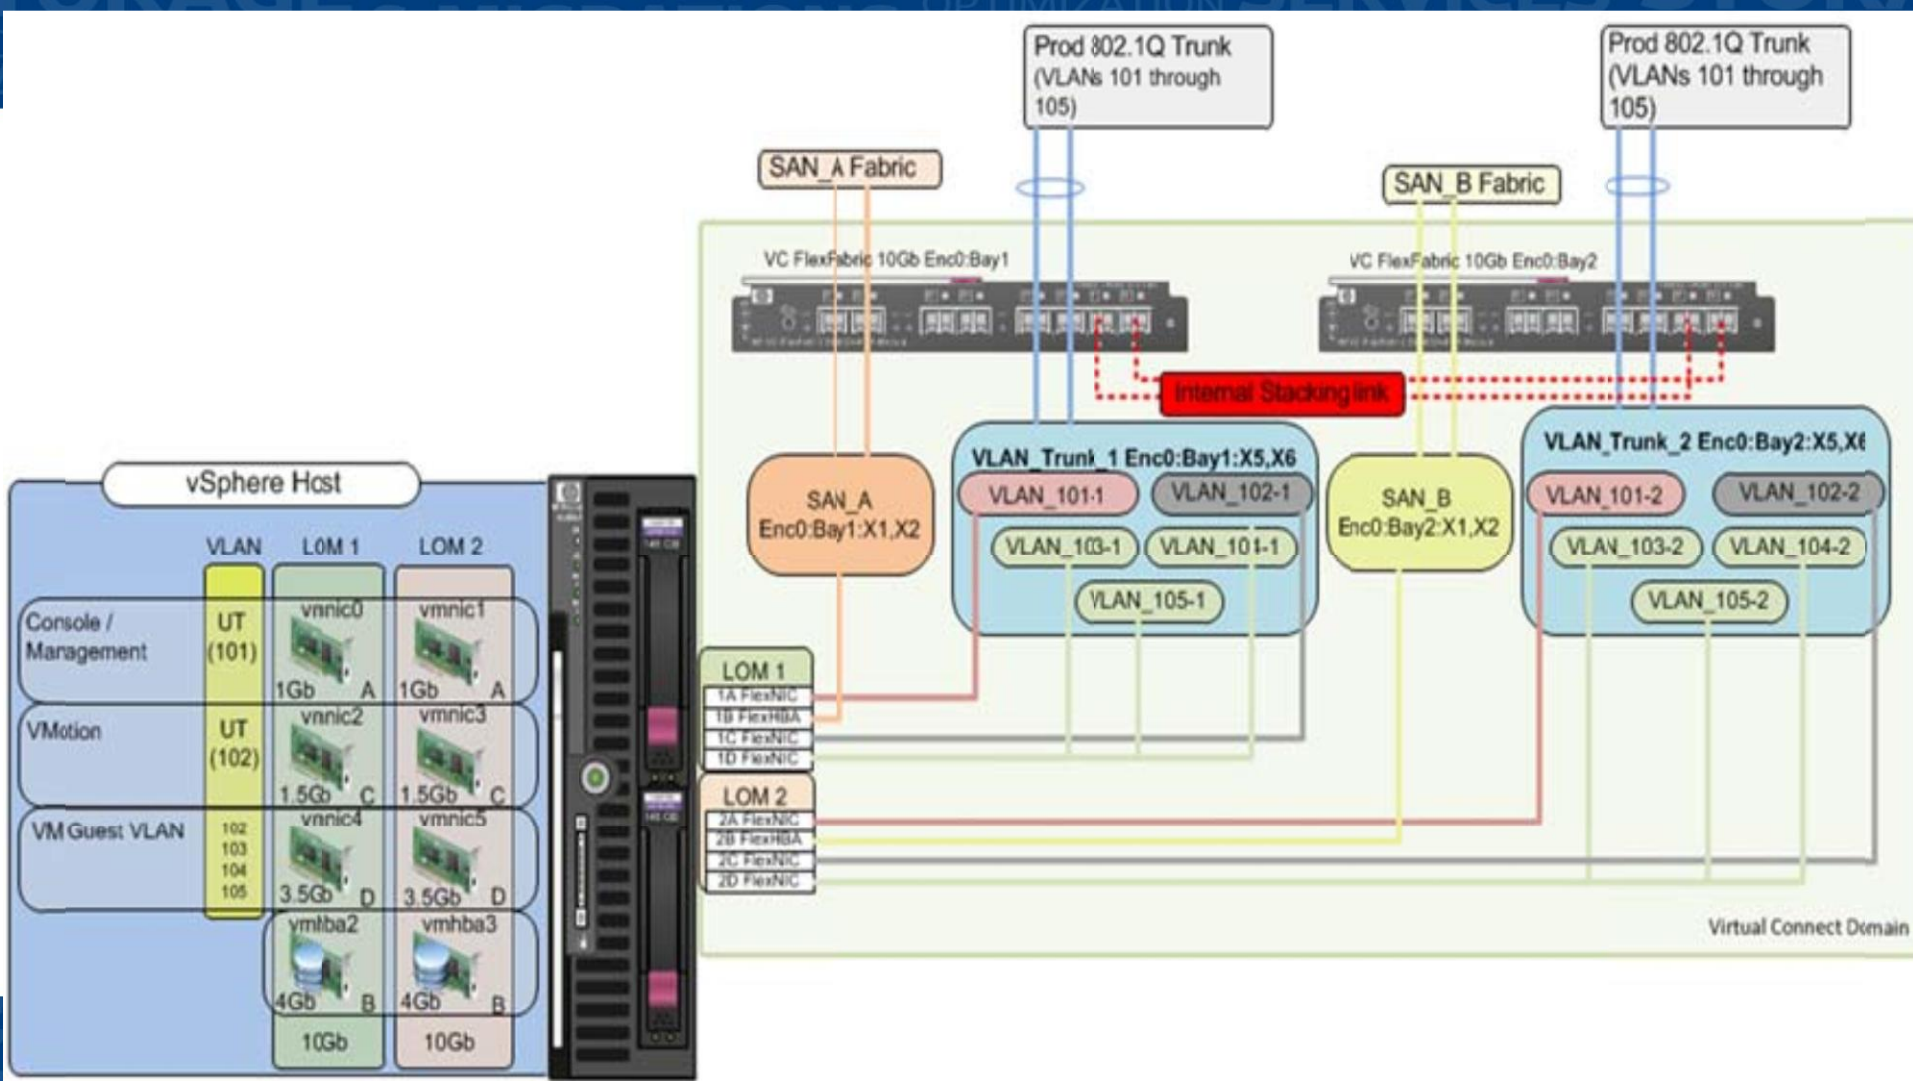
FCoE HBA Connections

Port	Connecte	FC SAN Name	Status	Port Speed	WWPN
1	Bay 1	Unassigned	?	DISABLED	VC-DEFINED
2	Bay 2	Unassigned	?	DISABLED	VC-DEFINED
3	Bay 3	Unassigned	?	DISABLED	VC-DEFINED
4	Bay 4	My_SAN_Fabric	✓	4	VC-DEFINED

+ Add

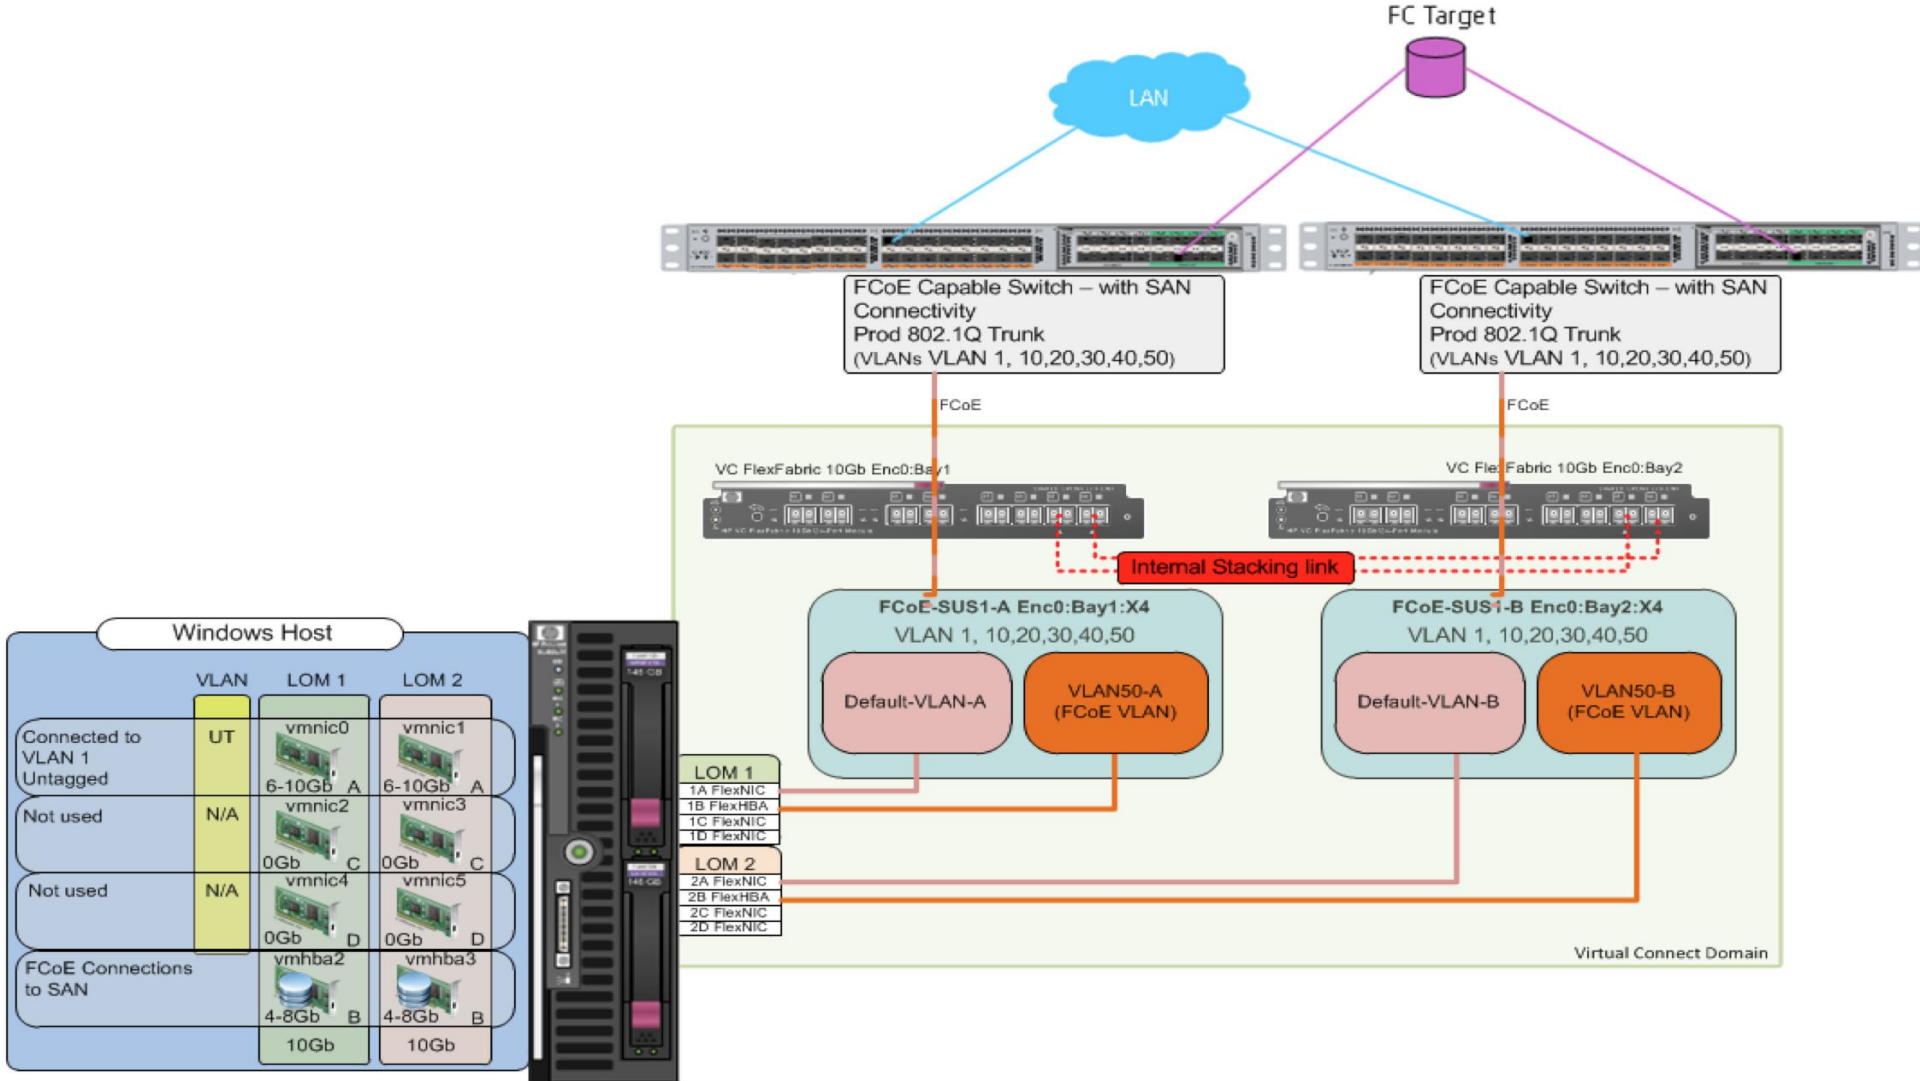
3. Ethernet design

- » Choice of Tunnel/Map – no longer necessary, wizard guides you
- » One VC Shared Uplink Set that contains all external ports per VC module – all ports become Active
 - » Why not go across modules? Because VC will put ports in 2nd module in Standby
- » Must enable
 - » SmartLink on each VCnet
 - » alternate between VC Shared Uplink Sets for each port in Profile,
 - » enable LACP PortChannels/Aggregation
 - » Virt.Port ID teaming/TLB on ESX



4. FCoE design

- » Converged Shared Uplink Sets (SUS) can contain both the FCoE network and non-FCoE networks
- » FCoE-capable SUS must always contain ports from a single VC module
- » Active/Active configuration for FCoE traffic is required
- » For Multi Enclosure (ME) environments, all corresponding ports in the remote enclosures will be included in the same SUS
 - » e.g. selecting enc0:bay1:X1 means bay1:X1 in all remote enclosures is also included.



In the server profile, alternate between Shared Uplink Sets for each NIC port. Watch out for a limitation on same Vlan to different LOMs (Eth 1, 3 – LOM1 –VLAN1 in red)

Edit Server Profile: Profile_FCoE1

Profile_VLAN1

Ethernet Adapter Connections

Port	Network Name	Status	Port Speed Type	Allocated Port Speed...	PXE	Multicast Filter	MAC	Mapping
1	VLAN1-B	✓	PREFERRED	4.9 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-02	LOM:1-a => Bay 1:d1.v1
2	VLAN1-B	✓	PREFERRED	6 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-04	LOM:2-a => Bay 2:d1.v1
3	VLAN1-B	✓	PREFERRED	4.9 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-06	LOM:1-c => Bay 1:d1.v3
4	VLAN10-B	✓	PREFERRED	2 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-08	LOM:2-c => Bay 2:d1.v3
5	VLAN10-B	✓	PREFERRED	100 M	USE-BIOS	None	00-17-A4-77-54-0A	LOM:1-d => Bay 1:d1.v4

+ Add

ISCSI HBA Connections

Port	Network Name	Status	Port Speed Type	Mapping
------	--------------	--------	-----------------	---------

+ Add

FCoE HBA Connections

Port	Connect	FC SAN / FCoE Network Name	Type	Status	MAC	Mapping	Act
1	Bay 1	FCoE-A	FCOE	✗	00-17-A4-77-54-00	LOM:1...	
2	Bay 2	FCoE-B	FCOE	✗	00-17-A4-77-54-01	LOM:2...	De

+ Add

Fibre Channel Boot Parameters

VCM Error

These connections are configured with **duplicate** networks. Ensure there are no duplicate networks on connections mapped to the same physical port.

Each of the following sets of connections is mapped to the same port:

Ethernet 1, 3 are mapped to LOM:1

NOTE: This dialog may be left open for reference while corrections are made.

OK

Same Vlan to different LOMs limitation avoided (uses VLAN1A for 1, 1B for 3).

Edit Server Profile: Profile_FCoE1

PROFILE_CULT | Details | [+]

Ethernet Adapter Connections

Port	Network Name	Status	Port Speed Type	Allocated Port Speed...	PXE	Multicast Filter	MAC	Mapping
1	VLAN1-A	✓	PREFERRED	4.9 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-02	LOM:1-a => Bay 1:d1.v1
2	VLAN1-B	✓	PREFERRED	6 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-04	LOM:2-a => Bay 2:d1.v1
3	VLAN1-B	✓	PREFERRED	4.9 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-06	LOM:1-c => Bay 1:d1.v3
4	VLAN10-B	✓	PREFERRED	2 Gb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-08	LOM:2-c => Bay 2:d1.v3
5	VLAN10-B	✓	PREFERRED	100 Mb - 10 Gb	USE-BIOS	None	00-17-A4-77-54-0A	LOM:1-d => Bay 1:d1.v4

+ Add

iSCSI HBA Connections

Port	Network Name	Status	Port Speed Type	Allocated Port Speed (Mi...	Boot Setting	MAC	Mapping
------	--------------	--------	-----------------	-----------------------------	--------------	-----	---------

+ Add

FCoE HBA Connections

4. Flex10 ESX 5.1 Tips

- » For a VMkernel interface enabled with Jumbo Frame with MTU of 9000 Bytes, ensure the Virtual Connect NIC is operating at a minimum of 1Gb/s.
- » For Fault Tolerance logging NICS, make sure you reserve at least 1 Gbit, separately from the internal VMotion network.
- » **#HPtrick** With rollback of settings, it is safe to use VMware's Distributed vSwitch for Management & Data VMs.

IV. EVA (P6000) Configuration for VMware

» Goals: Maximize
redundancy &
load balancing

HP EVA (P6000) and vSphere ESX 5.1

- In most cases, no longer need to load balance EVA manually – leave at default Multi-Path -- Round-Robin
- For Microsoft cluster RDM LUNs, use MRU policy
- Set host profile to VMware at the LUN level. If still using VMware VCB, then set VCB host profile to Win2003/2008.
- For LUN Configuration, alternate between controller A and controller B using the Path-A-Failover/Failback or Path-B-Failover/Failback setting
- Configure round robin advanced parameters to IOPS=1 for vSphere 4.x/5.
 - <http://www.yellow-bricks.com/2010/03/30/whats-the-point-of-setting-iops1/>

V. VMware ESX 5.1 (vSphere) Best Practices

» Goal: Optimize
VMware on Blades

1. Storage/Data Store tips

- » Establish a standard LUN size that you will use. 500 GB is a common choice. Now with VMFS5, 1TB or 2TB is good.
- » Use RDMs (Virtual) only for VM disks that exceed your standard LUN size. Only put data on RDMs, not the OS.
- » Thin provision in VMware and on the storage array for maximum efficiency, but setup alerts to not oversubscribe. Exceptions – Microsoft Exchange, very high perform. apps
- » Store ISOs and templates on a separate LUN
- » Use Paravirtualized adapters inside VMs for ALL I/O
- » **#HPTrick** Use max 10 heavy I/O, 15 medium, 20 light VMs on one LUN, **even** if you use HP EVA VAAI

VI. Summary

» Make sure you spend time on design before you build the solution. With c-Class Blades, Virtual Connect, vSphere, & HP EVA (P6000), planning is a must.

Thank you and Twitter chat

» I would like to thank Alex Kramer, Aboubacar Diare, Steve Mclean, Keenan Sugg, Doug Strain, Joanne McMenoman, and HP in general for documents & ideas, some of which I used here.

» **Chat with me on Twitter:**

June 18, 2013 -- Tuesday, 2-3 pm EST,

use hashtag: #HPtrick

Contact Info & Questions

Please fill out the Survey – it helps me come back.

Get the slides on my blog:

www.cdillc.com/newsroom/cloud-giraffe

Connect with me on LinkedIn: [Yury Magalif](#)

Twitter:

#HPtrick

@YuryMagalif

@CDILLC

YouTube:

www.youtube.com/user/ym640

Personal Blog: cloud-zebra.com

E-Mail Questions:

okzebra@gmail.com

yury.magalif@cdillc.com

