Deployment Guide

HP FlexFabric Reference Architecture – Data Center

Version 1

Table of Contents

Introduction

This document describes the implementation of data center (DC) architectures described in the FlexFabric Reference Architecture Guide. It outlines the deployment and technical details relevant to three HP FlexFabric models: Blade server one-tier, Top-of-Rack (ToR) Layer 2, and ToR Layer 3 designs.

Document structure

This document has 2 main sections:

1. FlexFabric design descriptions
    a. Offers a summary high level description of the solution type
2. Baseline configurations
    a. Describes the baseline configurations of each of these models

The document ends with the appendix section which provides a description of the primary technologies used in these types of deployments.

Audience

The document is intended for HP staff and deployment engineers that are building, or intend to build, a data center network and require design best practice recommendations and configuration examples.

Document Objectives

The document presents recommended designs for the DC network, and includes descriptions of a typical topology, software features, configuration guidelines, and other considerations relevant to an end to end design of a DC network using HP hardware. Since this solution utilizes standard based technologies, other manufacturers storage and server hardware may be substituted It is up to a designer engineer to choose the appropriate products for their network. All options shown provide configuration examples only.

FlexFabric designs

Blade server one-tier design

This type of network deployment signifies the current pinnacle of network virtualization. Server blades and enclosures connect to the DC core devices (Through Virtual Connect Flex-10/10D or FlexFabric modules), and allow for substantial compute density per rack, row, and data center. HP has optimized the BladeSystem server portfolio to support the vision and reality of virtualization. The network design optimizes the reality of high performance networking with simplicity. It allows flexibility in VM networking and converged I/O options. This approach is at the forefront of network design for virtualization, since it utilizes both the IRF framework and VC. It operates well while allowing for seamless management and troubleshooting of VMs.

Figure 1 Figure 1 Blade server one-tier design

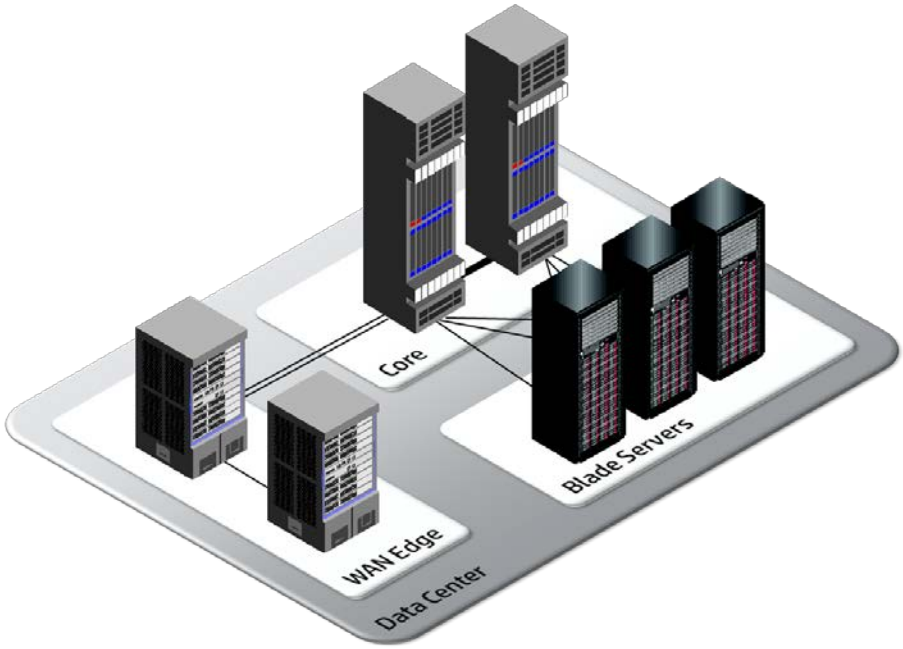

Benefits of a one-tier blade server design

The BladeSystem 1 Tier design provides the following benefits:

- Reduction of elimination of troublesome redundancy protocols such as Spanning-Tree and Virtual Router Redundancy Protocol (VRRP)
- Increase of usable bandwidth by allowing all links to be active as opposed to redundant links being blocked in a traditional design
- Faster convergence times in the event of a link, blade or unit failure. Up to 50 times faster than Rapid Spanning Tree
- Simpler design to implement and manage due to the reduction of the number of logical devices and related configuration files to be modified
- Decreased latency due to the elimination of aggregation layer.
- Virtual Connect Ethernet Converged modules offer:
    - LAN/SAN/Server configurations consolidated in a single H/W module using a unified management console with role based access control

- "Server Profiles" with virtualized MAC/WWN and NIC/HBAs, allowing servers to be easily added, configured, maintained, and moved
- Industry standard LAN/SAN protocols simultaneously fully supported (802.1q trunk, LACP, LLDP, FC, FCOE, H/W iSCSI and NFS)
- Direct storage attach option which eliminates the SAN purchase and reduce costs
- NIC virtualization into 4 separate networks each with their own bandwidth customizations

ToR Layer 2 / Layer 3 design

In a ToR design, servers connect to an access switch via copper or fiber within the rack, while the access switch connects to other consolidation or backbone switches within the data center.

Shorter copper Ethernet runs in the rack allow for multiple choices on cable types which can support various required speeds dictated by the systems in the rack. In many cases, server to switch connections can be up to 10GbE connections with support for integrated I/O. Longer runs from the rack to core usually utilize 10GbE MMF, however more bandwidth could be provided by using 40GbE or 100GbE ports. HP currently offers 40GbE ports on some HP switches, and will be expanding 40GbE coverage to more switches in the future. HP will also be offering 100GbE ports for even higher bandwidth links in the future.

Layer 2 ToR designs extend VLANs across the entire data center and are optimized for virtualization environments where VMs may be moving from rack to rack or data center to data center.

Layer 3 ToR designs inherently limit VM migration capabilities to a single rack but provide better scalability while isolating (containing) the effects of a failure locally.

Figure 2 ToR Layer 2 / Layer 3 design

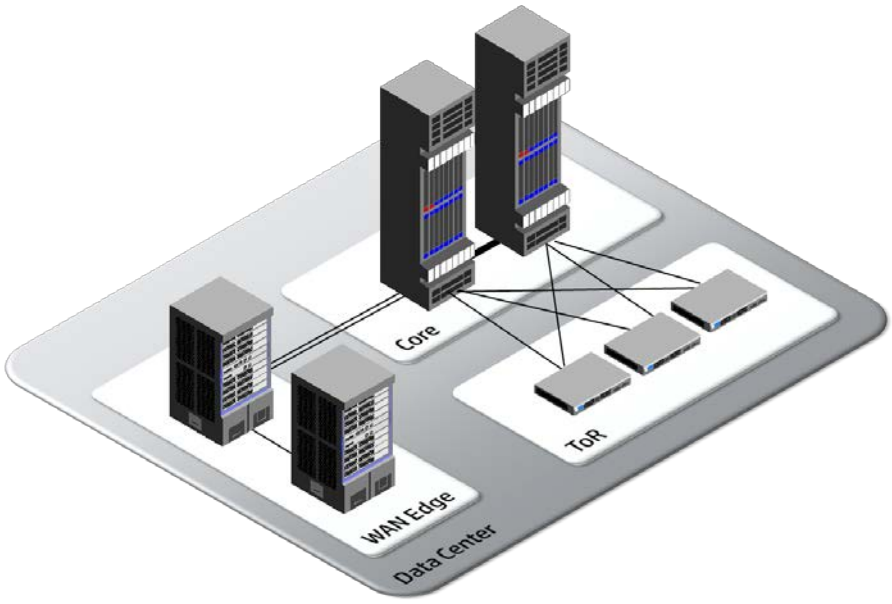

Baseline configurations

The following section will provide an end to end discussion on the 3 types of recommended data center FlexFabric deployments. Specifically, these will be FlexFabric 1-tier blade server, 2-tier L2 ToR, and 2-tier L3 ToR designs.

Note: Exact configurations listed will vary from specific customer implementations; however these configurations can be used as starting points and best practices and features discussed should still be considered relevant and important to follow.

Figure 3 Three data center designs

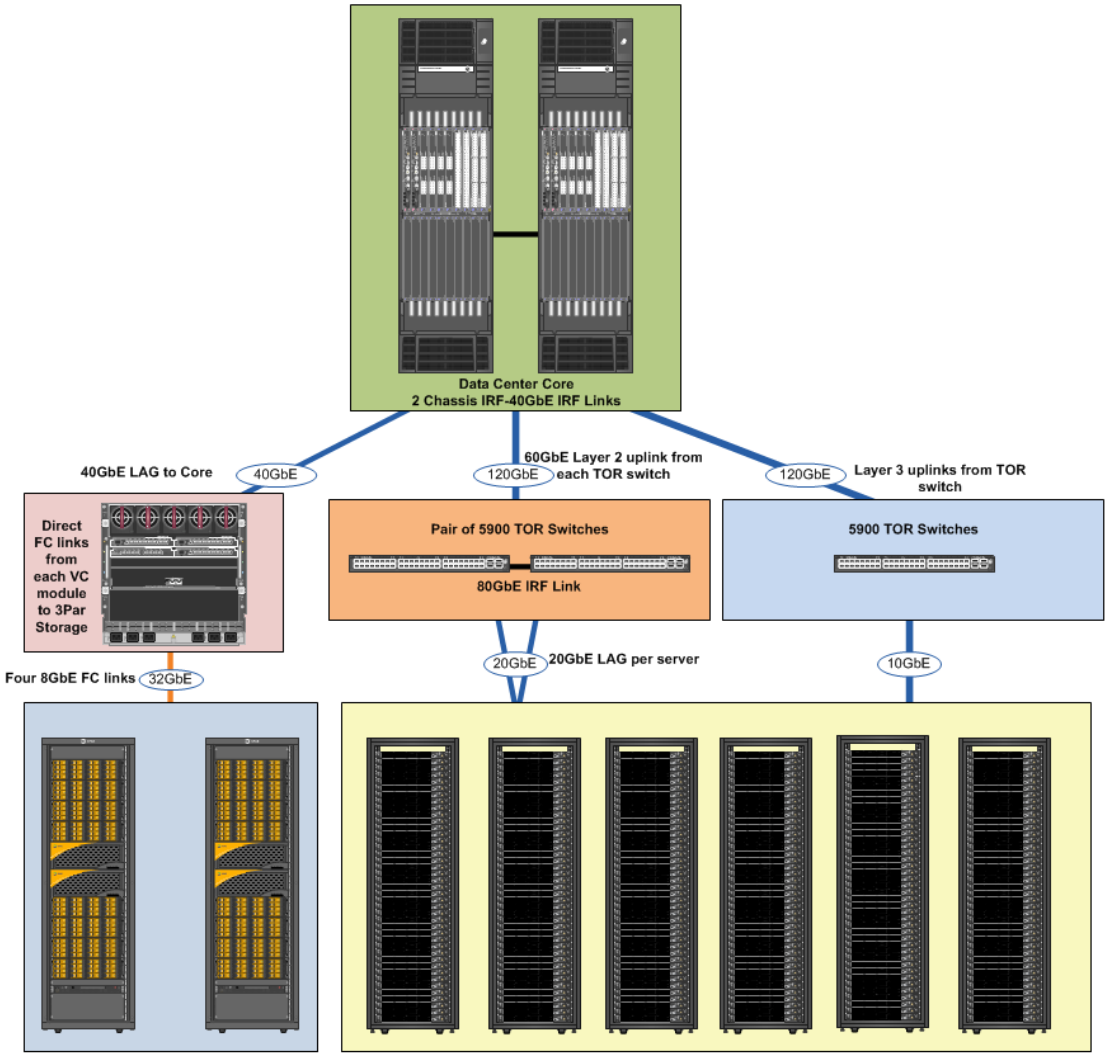

1-Tier Blade Server Design

The example 1-Tier Blade Server design will support both Fibre Channel and Ethernet. The design will utilize a pair of HP 12518 switches at the core which connect to a pair of HP Virtual Connect FlexFabric modules housed in an HP C7000 server enclosure with one or more server blades. The HP VC FlexFabric modules will then connect directly to a 3PAR FC storage array.

The pair of core switches will provide 20GbE (2*10GbE) connections to each FlexFabric module, providing a 40GbE active/active configuration. LACP will be used to aggregate those links, and multiple VLANs will be supported as the switch ports will be configured as trunk ports.

The solution will configure one HP server blade which will utilize 2 Virtual Connect FlexNICs (vNETS) to connect to the FlexFabric modules. The server NICs will each be configured to support 4Gb for Ethernet traffic, 2Gb for internal vMotion traffic, and 4Gb for Fibre Channel traffic (The configuration for one server in the enclosure is shown, however the configuration for additional enclosures would be similar). The HP Virtual Connect FlexFabric modules will utilize a pair of Shared Uplink Sets (SUS) which provide support for carrying multiple VLANs.

Additionally, two 8Gb FC uplinks from each FlexFabric module will DirectAttach connect to an HP 3PAR storage array for a complete end to end DC design.

From a physical view:

- The 12518 switches will utilize 4 10GbE links for IRF (ports used listed below)
- The pair of 12518s will utilize a 20GbE LAG which connects to the HP Virtual Connect FlexFabric module in bay 1 (ports used listed below)
- The pair of 12518s will utilize a 20GbE LAG which connects to the HP Virtual Connect FlexFabric module in bay 2 (ports used listed below)
- Each HP Virtual Connect FlexFabric module will use 2 8Gb FC links to DirectAttach to a 3PAR storage array (ports used listed below)

The scalability of this design is limited by the following aspects:

- As of this writing the 12518 can support a maximum of 288 line rate 10GbE ports or 576 10GbE ports at a 4:1 oversubscription ratio
- Using LEF modules, the 12518 supports 256K MAC address and 64K ARP entries as well as 1M IPv4 FIB entries
- The 12500 series of switches can support 4 chassis IRF for a total of 1,152 line rate 10GbE ports or 2,304 10GbE ports at a 4:1 oversubscription ratio
- The maximum number of Link Aggregation Groups (LAGs) is 240 with 12 ports per LAG
- The Virtual Connect FlexFabric module supports 32K MAC address entries
- The Virtual Connect FlexFabric module supports 4,096 networks in tunneled VLAN and 1,000 VLANs per Virtual Connect Ethernet domain
- Each physical FlexNIC can support up to 162 VLANs

Taking in the above criteria, an IRF pair of 12518s with line rate 10GbE ports can accommodate:

- 2,256 physical servers at a 4:1 over-subscription ratio
    - 288 line rate 10GbE ports per chassis *2=576
    - Minus 4 ports per chassis for IRF=568
    - Minus 2 port per chassis for WAN links=564
    - 4 link to each enclosure 566/4=141
    - 141 enclosures * 16 blades per=2,256

Four 12518 with line rate 10GbE ports can accommodate:

- 4,464 physical servers at 4:1 oversubscription
    - 288 line rate 10GbE ports per chassis *4=1,152
    - Minus 32 ports for IRF=1,120
    - Minus 1 port per chassis for WAN links=1,116
    - 4 links to each enclosure 1,116/4=279
    - 279*16=4,464

Notes: The maximum number of LAG groups for the 12518 is 240. If we subtract one for the WAN links the maximum number of supported enclosures is 239 even though the 10GbE density could support 279 enclosures.

The 1-Tier Blade Server example solution below utilizes 2 LAG groups per enclosure for redundancy. When using 2 LAG groups per enclosure the maximum number of supported enclosures is 119.

When scaling this design for virtualized environments, please note that the least common denominator is the blade enclosure interconnect module that supports 32K MAC entries. With 2 chassis IRF each server could support up to 14 VMs while a 4 chassis design would be limited to 8 VMs per server.

Figure 4 1-Tier Blade Server Physical View

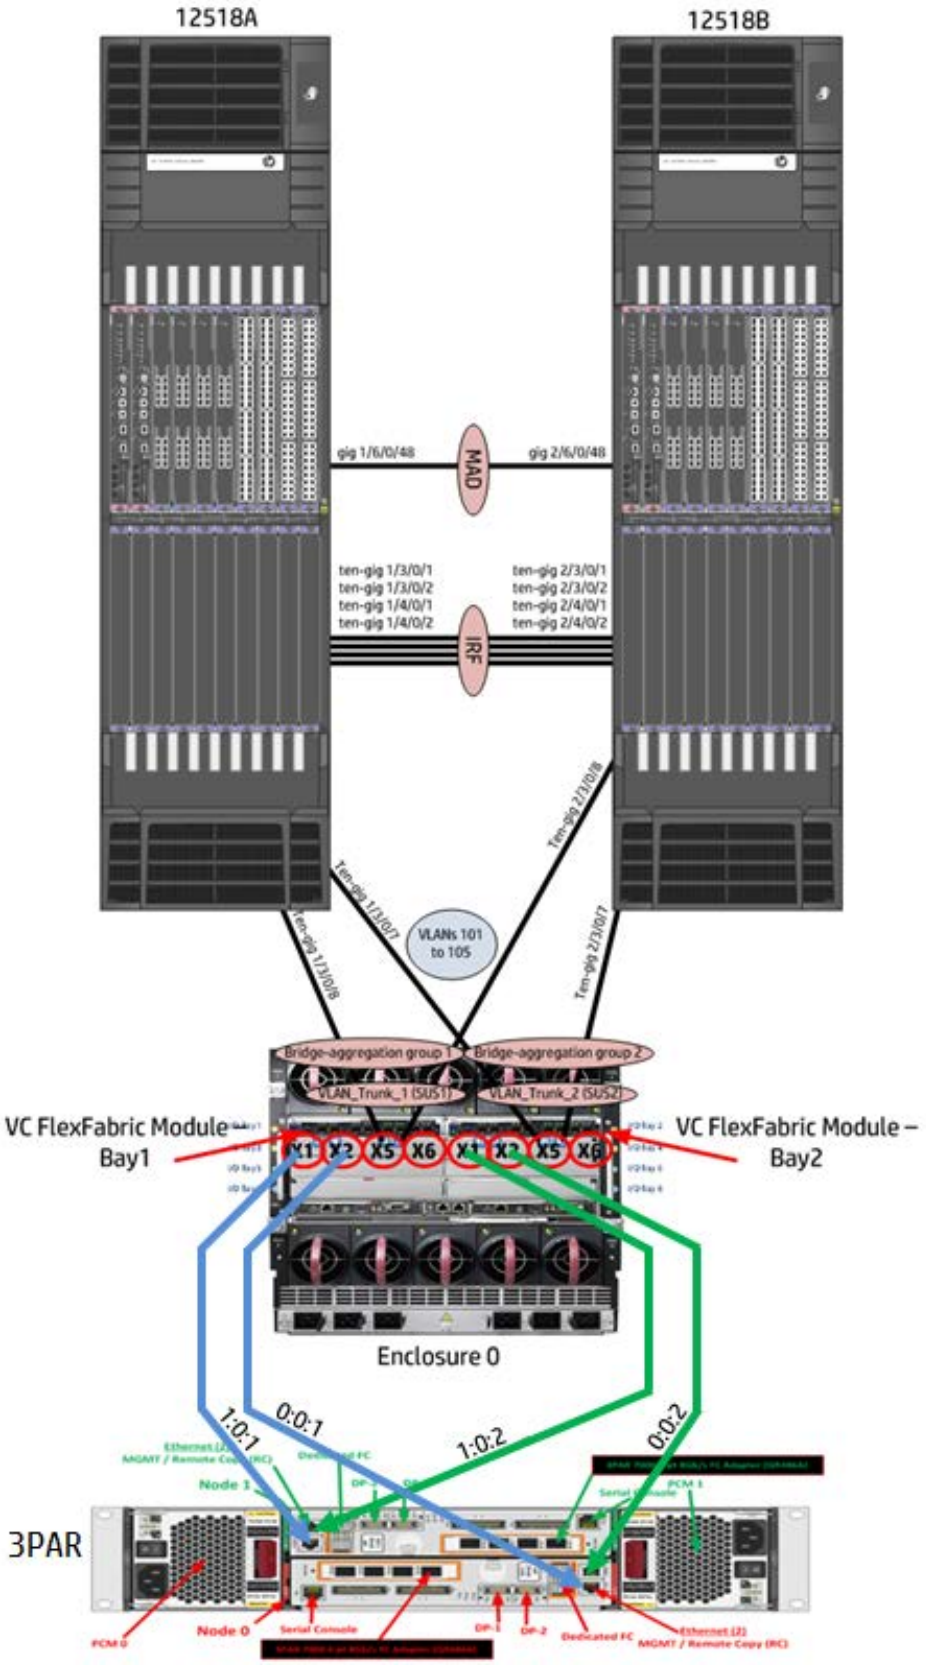

Table 1 1-Tier Blade Server Cabling Information

Link Type	Device	Port	Link Type	Remote Device	Port
IRF	12518A	1/3/0/1	10GbE	12518B	2/3/0/1
		1/3/0/2			2/3/0/2
		1/4/0/1			2/4/0/1
		1/4/0/2			2/4/0/2
IRF MAD	12518A	1/6/0/48	1GbE	12518B	2/6/0/48
To VC2	12518A	1/3/0/7	10GbE	VC FlexFabric – Bay 2	X5
To VC1	12518A	1/3/0/8	10GbE	VC FlexFabric – Bay 1	X5
To VC2	12518B	2/3/0/7	10GbE	VC FlexFabric – Bay 2	X6
To VC1	12518B	2/3/0/8	10GbE	VC FlexFabric – Bay 1	X6
To SAN A	VC FlexFabric – Bay 1	X1	8Gb FC	3PAR	1:0:1
		X2			0:0:1
To SAN B	VC FlexFabric – Bay 2	X1	8Gb FC		1:0:2
		X2			0:0:2

From a logical view:

- The server blade profile will be configured with two FlexNICs (vNet) and two FCoE FlexHBAs
- NICs 1 and 2 are connected to VLAN-101-x which is part of the Shared Uplink Sets, VLAN-Trunk-1 and VLAN-Trunk-2, respectively
- VLAN 101 will be configured as a native untagged VLAN, while VLANs 102, 103, 104, and 105 will be configured to pass through the VC modules to the server as tagged VLANs
- The VLAN-Trunks are connected, at 40Gb, to the 12518 core switches (2 pairs of 20GbE LACP trunks = 40GbE total bandwidth)
- The first FlexNIC (vNet) on each server LOM port has been configured to support up to 4Gb each
- The second FlexNIC (vNet) FCoE port on each server LOM port has been configured to support up to 4Gb each
- The third FlexNIC (vNet) on each server LOM port has been configured to support up to 2Gb for internal vMotion

Figure 5 Logical view of VC

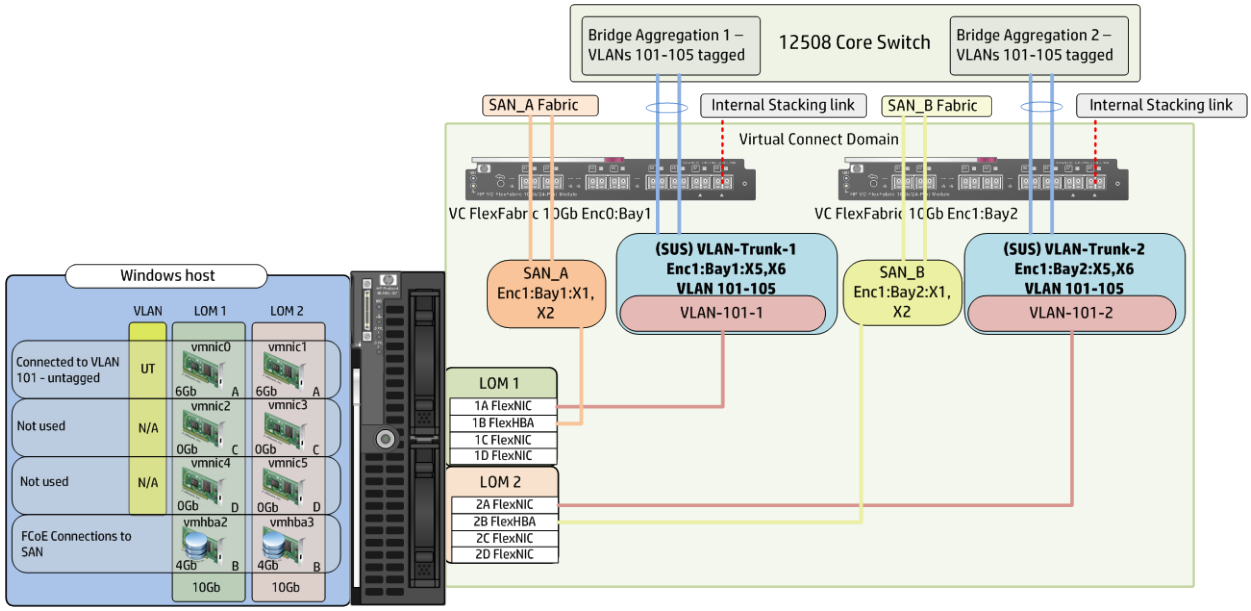

Table 2 1-Tier Blade Server VLANs and IP Addresses

Device	VLAN	IP Address	Uplink	Details
12518A (and B)	VLAN101	10.10.101.1 / 24	Bridge-aggregation group 1 & 2	Trunk (tagged)
	VLAN102	10.10.102.1 / 24		
	VLAN103	10.10.103.1 / 24		
	VLAN104	10.10.104.1 / 24		
	VLAN105	10.10.105.1 / 24		
12518A	VLAN3	192.168.3.1 / 24	1/6/0/48	MAD IRF member 1
12518B	VLAN3	192.168.3.2 / 24	2/6/0/48	MAD IRF member 2
VC FF – Bay 1	VLAN101	N/A	Shared Uplink Set (VLAN-Trunk-1)	Native
	VLAN102			Tagged
	VLAN103			
	VLAN104			
	VLAN105			
VC FF – Bay 2	VLAN101	N/A	Shared Uplink Set (VLAN-Trunk-2)	Native
	VLAN102			Tagged
	VLAN103			
	VLAN104			
	VLAN105			

Table 3 LAGs and SUSs

LAG	Local Ports	Remote Device	SUS
Bridge-aggregation group 1	1/3/0/8	VC FF Bay 1 – X5	VLAN-Trunk-1
	2/3/0/8	VC FF Bay 1 – X6	
Bridge-aggregation group 2	1/3/0/7	VC FF Bay 2 – X5	VLAN-Trunk-2
	2/3/0/7	VC FF Bay 2 – X6	

LAN Deployment Procedure

This section describes the configuration steps to deploy the HP 12518 platforms using IRF for the 1-Tier blade server design.

- Rack the chassis and install all necessary physical components to establish a functional HP 12518
- Configure IRF
  - Configure IRF
  - Configure MAD
- Configure VLANs
  - Create VLAN, names, descriptions
  - Configure VLAN interface IP addresses
  - Undo shutdown VLAN interfaces
  - Configure VLAN MTU, if necessary
- Create and establish uplinks and downlinks
  - Configure Link Aggregation Control Protocol (LACP) for each set of uplinks and downlinks
  - Treat ports attached to Virtual Connect modules as edge port
    - Enable STP edge port and BPDU guard
  - Configure port descriptions to assist in troubleshooting
  - Configure load sharing

LAN Configuration Procedures

Rack the chassis and install all necessary physical components to establish a functional HP 12518.

- Detailed descriptions on how to install the HP 12518 switches is out of scope of this document, however a chart of the high level steps that should be followed is listed below. Make sure to fully review the HP 12500 installation manual to ensure proper installation
- http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?contentType=SupportManual&lang=en&cc=us&docIndexId=64179&taskId=101&prodTypeId=12883&prodSeriesId=4177453

Figure 6 12500 installation flow chart

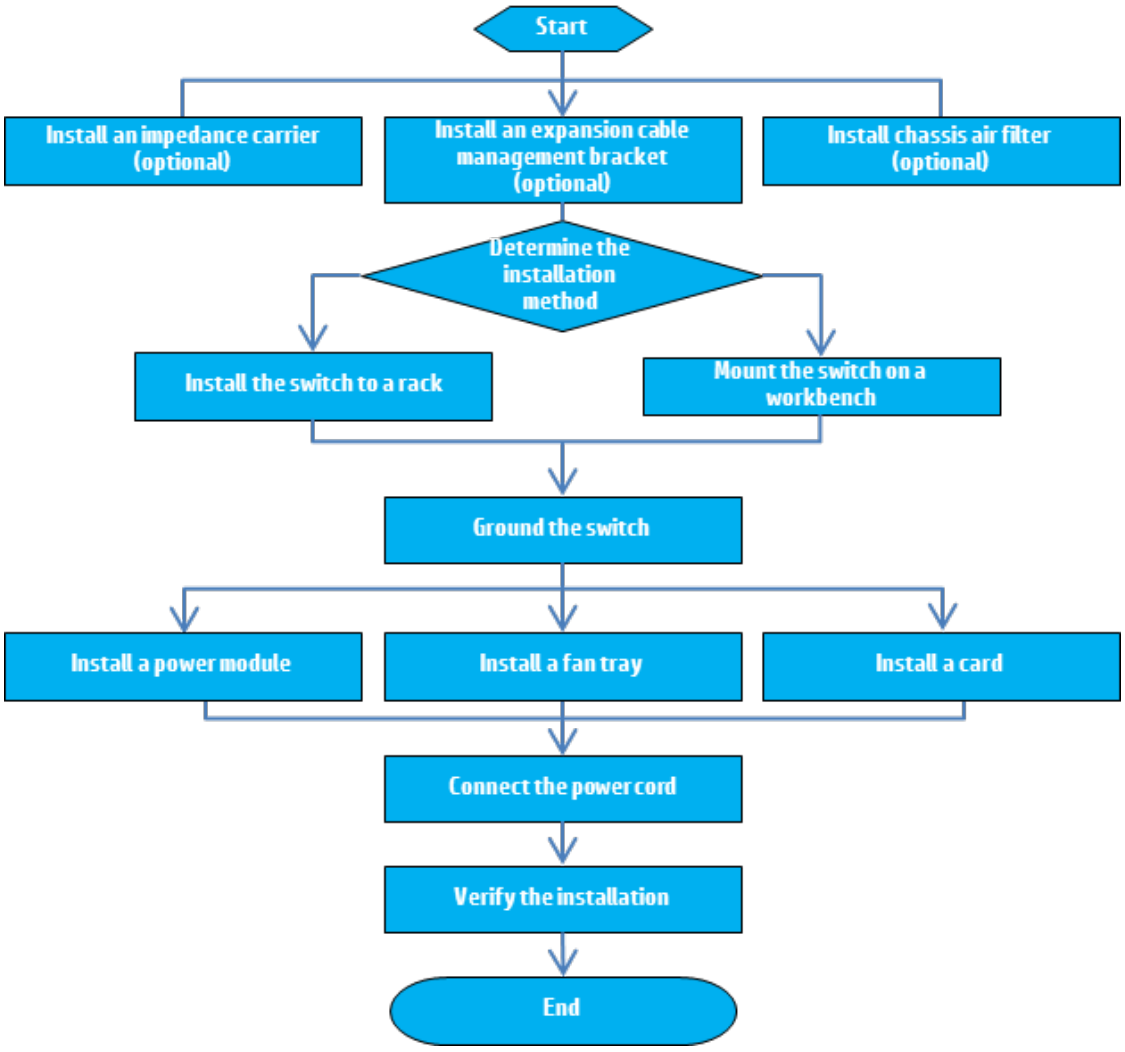

Configure IRF & BFD MAD detection

- The configuration listed refers to the two HP 12518 switches as 12518A and 12518B. Once IRF configuration has been established the same will refer to the master which will be 12518A
- To offset the risk of IRF virtual device partition, configure MAD to detect multi-active collisions. In this example, BFD MAD is adopted

Configure 12518A

1) Set the member ID of 12518A to 1, change the IRF priority to 30, then save the configuration. The larger IRF priority will determine the IRF Master

```
irf member 1
irf priority 30
save
```

2) Change the operating mode to IRF. The switch automatically reboots to make the change effective

```
chassis convert mode irf
```

3) Shutdown the IRF interfaces

```
interface ten-gigabitethernet 1/3/0/1
  shutdown
interface ten-gigabitethernet 1/3/0/2
  shutdown
interface ten-gigabitethernet 1/4/0/1
  shutdown
interface ten-gigabitethernet 1/4/0/2
  shutdown
  quit
```

4) Create IRF port 1/2, and bind the physical IRF ports to it

```
irf-port 1/2
  port group interface ten-gigabitethernet 1/3/0/1
  port group interface ten-gigabitethernet 1/3/0/2
  port group interface ten-gigabitethernet 1/4/0/1
  port group interface ten-gigabitethernet 1/4/0/2
  quit
```

5) Undo the Shutdown on the IRF interfaces and save the configuration

```
interface ten-gigabitethernet 1/3/0/1
  undo shutdown
interface ten-gigabitethernet 1/3/0/2
  undo shutdown
interface ten-gigabitethernet 1/4/0/1
  undo shutdown
interface ten-gigabitethernet 1/4/0/2
  undo shutdown
  quit
save
The current configuration will be written to the device. Are you
sure? [Y/N]:y
Please input the file name(*.cfg)[flash:/startup.cfg]
(To leave the existing filename unchanged, press the enter key):
flash:/startup.cfg exists, overwrite? [Y/N]:y
Validating file. Please wait.....................................
The current configuration is saved to the active main board
successfully.
Configuration is saved to device successfully.
```

Configure 12518B

1) Set the member ID of 12518B to 2 and specify the priority. If this switch is going to be an IRF slave, there is no need to modify the default priority of 1. Save the configuration

```
irf member 2
save
```

2) Change the operating mode to IRF. The switch automatically reboots to make the change effective

```
chassis convert mode irf
```

3) Shutdown the IRF interfaces

```
interface ten-gigabitethernet 2/3/0/1
  shutdown
interface ten-gigabitethernet 2/3/0/2
  shutdown
interface ten-gigabitethernet 2/4/0/1
  shutdown
interface ten-gigabitethernet 2/4/0/2
  shutdown
  quit
```

4) Create IRF port 1/1, and bind the physical IRF ports to it

```
irf-port 1/1
  port group interface ten-gigabitethernet 2/3/0/1
  port group interface ten-gigabitethernet 2/3/0/2
  port group interface ten-gigabitethernet 2/4/0/1
  port group interface ten-gigabitethernet 2/4/0/2
  quit
```

5) Undo the Shutdown on the IRF interfaces and save the configuration

```
interface ten-gigabitethernet 2/3/0/1
  undo shutdown
interface ten-gigabitethernet 2/3/0/2
  undo shutdown
interface ten-gigabitethernet 2/4/0/1
  undo shutdown
interface ten-gigabitethernet 2/4/0/2
  undo shutdown
  quit
save
The current configuration will be written to the device. Are you
sure? [Y/N]:y
Please input the file name(*.cfg)[flash:/startup.cfg]
(To leave the existing filename unchanged, press the enter key):
flash:/startup.cfg exists, overwrite? [Y/N]:y
Validating file. Please wait.....................................
The current configuration is saved to the active main board
successfully.
Configuration is saved to device successfully.
```

Connect the two switches with the appropriate cables and bring up the IRF fabric

1) Turn off both switches
2) Connect the IRF links
3) Turn on the first chassis and wait until it finishes the boot cycle
4) Turn on the second chassis
5) When both switches have completed their boot cycles, verify the configuration:

```
display irf
Switch Role Priority CPU-Mac Description
*+1 Master 30 b8af-6731-c7ad -----
2 Slave 1 b8af-672c-4308 -----
-------------------------------------------------
  * indicates the device is the master.
  + indicates the device through which the user logs in.
  The Bridge MAC of the IRF is: b8af-6731-c77a
  Auto upgrade          : yes
  Mac persistent        : always
  Domain ID             : 1

display irf configuration
MemberID    NewID       IRF-Port1         IRF-Port2
1     1           disable           Ten-Gig1/3/0/1
Ten-Gig1/3/0/2
Ten-Gig1/4/0/1
Ten-Gig1/4/0/2
2     2           Ten-Gig2/3/0/1    disable
Ten-Gig2/3/0/2
Ten-Gig2/4/0/1
Ten-Gig2/4/0/2
```

Configure BFD MAD in case of IRF split

1) Create VLAN 3, and add required ports to VLAN 3. Note that for redundancy purposes, a LAG group could be created that contains more than one link. This example shows one link.

```
vlan 3
  port gigabitethernet 1/6/0/48 gigabitethernet 2/6/0/48
  quit
```

2) Create VLAN-interface 3 and configure a MAD IP address for the interface

```
interface vlan-interface 3
  mad bfd enable
  mad ip address 192.168.3.1 24 member 1
  mad ip address 192.168.3.2 24 member 2
  quit
```

3) Because BFD MAD detection and spanning tree function are mutually exclusive, disable the spanning tree function on the MAD ports.

```
interface Gigabitethernet 1/6/0/48
  undo stp enable
quit
  interface Gigabitethernet 2/6/0/48
  undo stp enable
```

Create VLAN-interface and configure its IP address.

1) Create VLAN-interface 101 and configure its IP address as 10.10.101.1/24

```
vlan 101
```

```
  description To Virtual Connect
  quit
interface vlan-interface 101
  ip address 10.10.101.1 24
  undo shutdown
  quit
```

2) Create VLAN-interface 102 and configure its IP address as 10.10.102.1/24

```
vlan 102
  description To Virtual Connect
  quit
interface vlan-interface 102
  ip address 10.10.102.1 24
  undo shutdown
  quit
```

3) Create VLAN-interface 103 and configure its IP address as 10.10.103.1/24

```
vlan 103
  description To Virtual Connect
  quit
interface vlan-interface 103
  ip address 10.10.103.1 24
  undo shutdown
  quit
```

4) Create VLAN-interface 104 and configure its IP address as 10.10.104.1/24

```
vlan 104
  description To Virtual Connect
  quit
interface vlan-interface 104
  ip address 10.10.104.1 24
  undo shutdown
quit
```

5) Create VLAN-interface 105 and configure its IP address as 10.10.105.1/24

```
vlan 105
  description To Virtual Connect
  quit
interface vlan-interface 105
  ip address 10.10.105.1 24
  undo shutdown
  quit
```

Link Aggregation Configuration Procedure

- Configure 2 dynamic Link-Aggregation groups (1 and 2) and assign appropriate ports
- Configure ports as stp edged ports and enable BPDU protection
- Configure each LAG as trunk ports and ensure that they each carry tagged VLANs 101 to 105

1) Create Layer 2 aggregate interface Bridge-Aggregation 1, and configure the link aggregation mode as dynamic

```
interface bridge-aggregation 1
  link-aggregation mode dynamic
  quit
```

2) Configure Layer 2 aggregate interface Bridge-Aggregation 1 as a trunk port and assign it to VLANs 101 to 105

```
interface bridge-aggregation 1
  port link-type trunk
  port trunk permit vlan 101 to 105
  quit
```

3) Assign ports to link aggregation group 1 one at a time, undo shutdown, and configure as stp edge ports

```
interface Ten-GigabitEthernet 1/3/0/8
  description To VC 1
  port link-aggregation group 1
  stp edged-port
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/3/0/8
  description To VC 1
  port link-aggregation group 1
  stp edged-port
  undo shutdown
  quit
```

4) Create Layer 2 aggregate interface Bridge-Aggregation 2, and configure the link aggregation mode as dynamic

```
interface bridge-aggregation 2
  link-aggregation mode dynamic
  quit
```

5) Configure Layer 2 aggregate interface Bridge-Aggregation 2 as a trunk port and assign it to VLANs 101 to 105

```
interface bridge-aggregation 2
  port link-type trunk
  port trunk permit vlan 101 to 105
  quit
```

6) Assign ports to link aggregation group 2 one at a time, undo shutdown, and configure as stp edge ports

```
interface Ten-GigabitEthernet 1/3/0/7
  description To VC 2
  port link-aggregation group 2
  stp edged-port
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/3/0/7
  description To VC 1
```

```
port link-aggregation group 2
stp edged-port
undo shutdown
quit
```

7) Enable BPDU guard so that when edge ports receive configuration BPDUs, the system will close these ports and notify the NMS that these ports have been closed by the spanning tree protocol

```
stp bpdu-protection
```

8) Configure the 12518A to use either the source and destination IP or the source and destination MAC addresses as the global link-aggregation load sharing criteria. In virtual environments, consider using source/destination IP load-sharing as MAC load-sharing may favor a single link. Both options are shown below.

```
link-aggregation load-sharing mode source-ip destination-ip
link-aggregation load-sharing mode source-mac destination-mac
```

Virtual Connect Deployment Procedure

The following Virtual Connect deployment procedures and configurations assume that the VC modules are already installed into an existing C3000/7000 enclosure. The configuration sections provide instructions on setting up the primary features required for connectivity to the LAN and SAN.

- Physically connect interface Ten-GigabitEthernet 1/3/0/7 on the 12518 switches to port X5 of the VC module in Bay 2
- Physically connect interface Ten-GigabitEthernet 2/3/0/7 on the 12518 switches to port X6 of the VC module in Bay 2
- Physically connect interface Ten-GigabitEthernet 1/3/0/8 on the 12518 switches to port X5 of the VC module in Bay 1
- Physically connect interface Ten-GigabitEthernet 2/3/0/8 on the 12518 switches to port X6 of the VC module in Bay 1
- Physically connect Ports X1 and X2 on the FlexFabric module in Bay 1 to ports 1:0:1 and 0:0:1 in the 3PAR 7200
- Physically connect Ports X1 and X2 on the FlexFabric module in Bay 2 to ports 1:0:2 and 0:0:2 in the 3PAR 7200
- Define a new Shared Uplink Set (VLAN-Trunk-1)
- Define a new Shared Uplink Set (VLAN-Trunk-2) (Copying a Shared Uplink Set)
- Add vMotion network
- Define a new FC SAN Fabric
- Define a Server Profile

Notes: Complete instructions on setting up VC modules into new enclosures is out of scope for this document. The configurations provide an introduction to VC modules and are basic configurations designed to provide general connectivity. Refer to the HP Virtual Connect Technology web page for access to complete information on VC modules and configurations: http://h18004.www1.hp.com/products/blades/virtualconnect/index.html

Direct Attach to 3PAR support is only available for VC FlexFabric module running VC3.70 and later.

3PAR Port Persistence is not supported when using a Virtual Connect Direct Attach to 3PAR solution. In the event of a node failure within the 3PAR 7200, path failover will occur, however depending on the actual configuration, there will be a I/O pause of roughly 10 seconds.

Virtual Connect Configuration Procedures

Figure 7 Virtual Connect Manager home screen

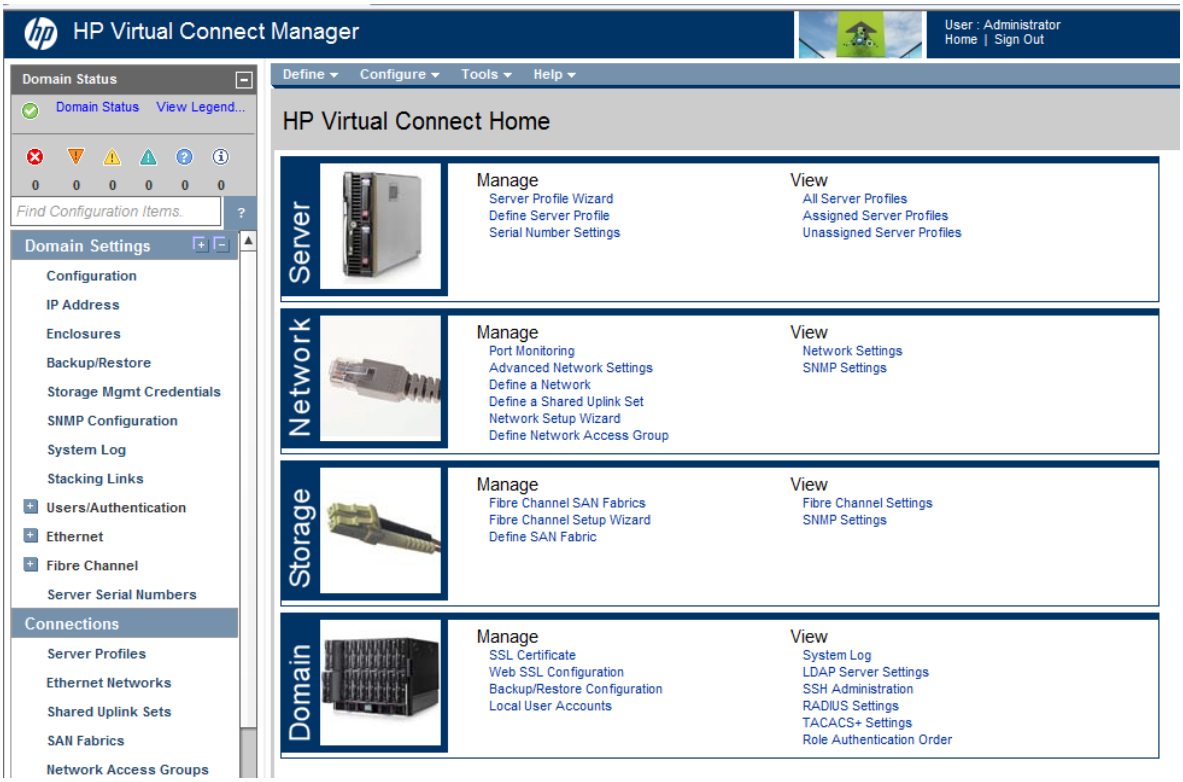

It is recommended that you utilize the Virtual Connect Ethernet Manager GUI to configure the Virtual Connect Modules. However, many of the configuration settings within VC can also be accomplished via a CLI command set.

1) Connect the physical ports as shown under the design details section, then log into the GUI of the onboard Run the Virtual Connect Manager (VCM) Domain Setup Wizard

Figure 8: VC Access via Onboard Administrator

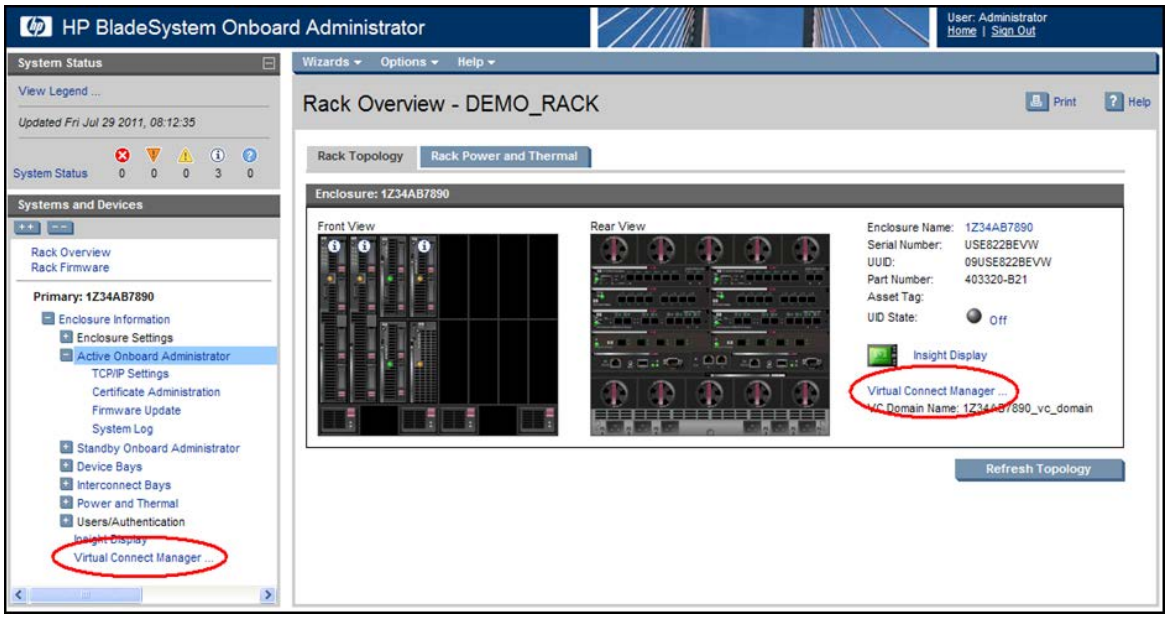

2) Create a SUS named VLAN-Trunk-1 and connect it to FlexFabric Ports X5 and X6 on Module 1. These ports will connect to the IRF'ed 12518 switches as shown in the 1-Tier Blade Server design section

```
On the Virtual Connect Home page select “Define”, and then
“Shared Uplink Set”.
Insert Uplink Set Name as “VLAN-Trunk-1”.
Select Add Port, then add the following port;
Enclosure 1, Bay 1, Port X5
Enclosure 1, Bay 1, Port X6
```

Figure 9 Define Shared Uplink Set

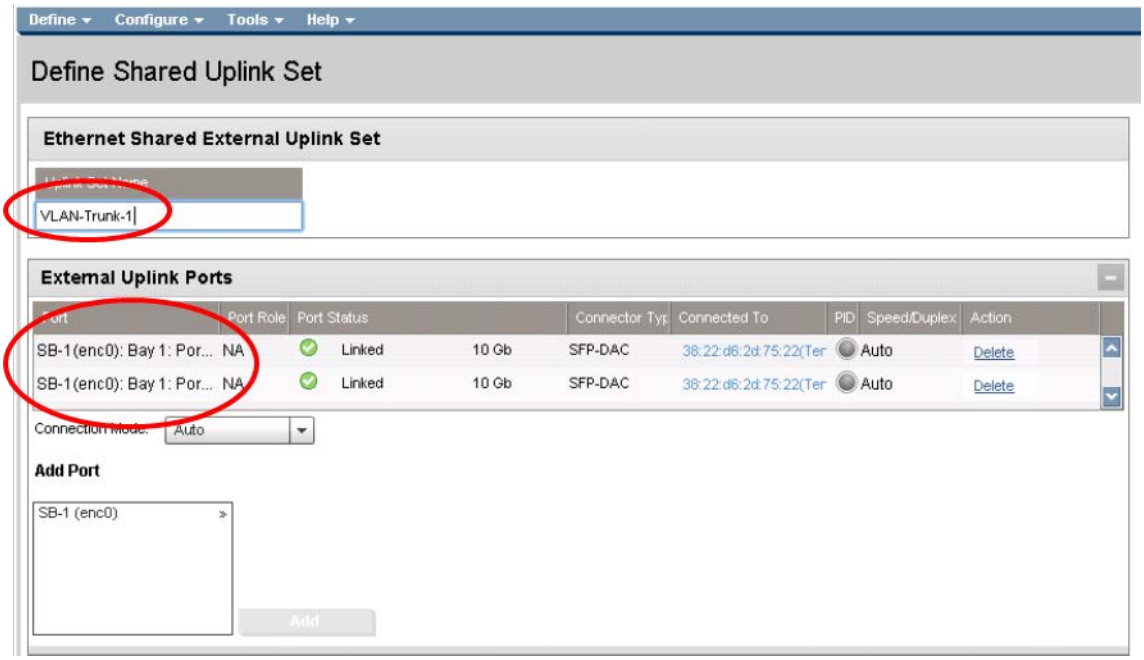

3) Associate networks with the SUS, enable SmartLink, and apply

Click Add under Associated Networks and select the Multiple Networks radio button and add the following VLANs;
Enter Name as VLAN-
Enter Suffix as -1
Enter VLAN IDs as follows (and shown in the graphic below;
101-105
Enable SmartLink on ALL networks
Click Apply
After clicking apply, a list of VLANs will be presented as shown below.

Figure 10 Defining SUS – associating networks

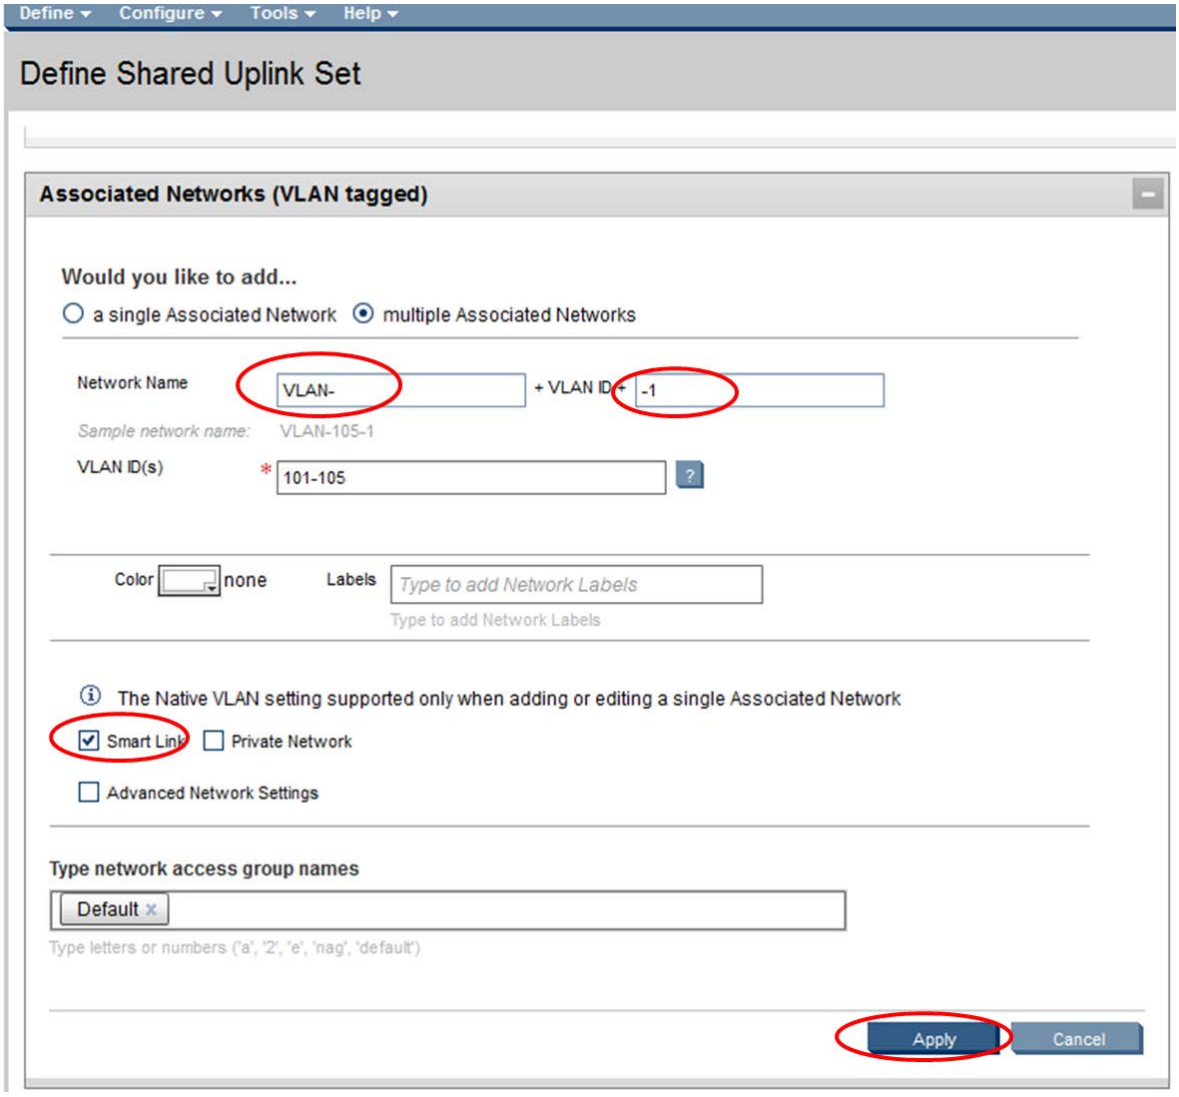

Note: You can optionally specify a network "color" or "Label" when creating a shared Uplinkset and its networks. In the example above we have not set either color or label.

Figure 11 List of VLAN created after defining VLAN-Trunk-1 SUS

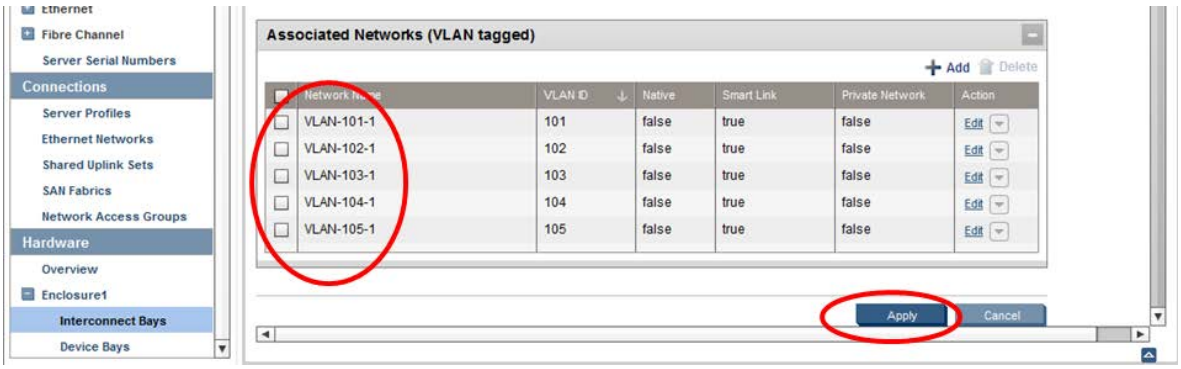

4) Configure native VLAN, and apply

```
To configure one of the VLANs, in this case VLAN101, as a Native
or Untagged VLAN
Click on Edit next to VLAN-101-1
Click on Native
Click on Apply
Click on Apply
```

Figure 12 Configuring VLAN-101-1 as Native/default VLAN

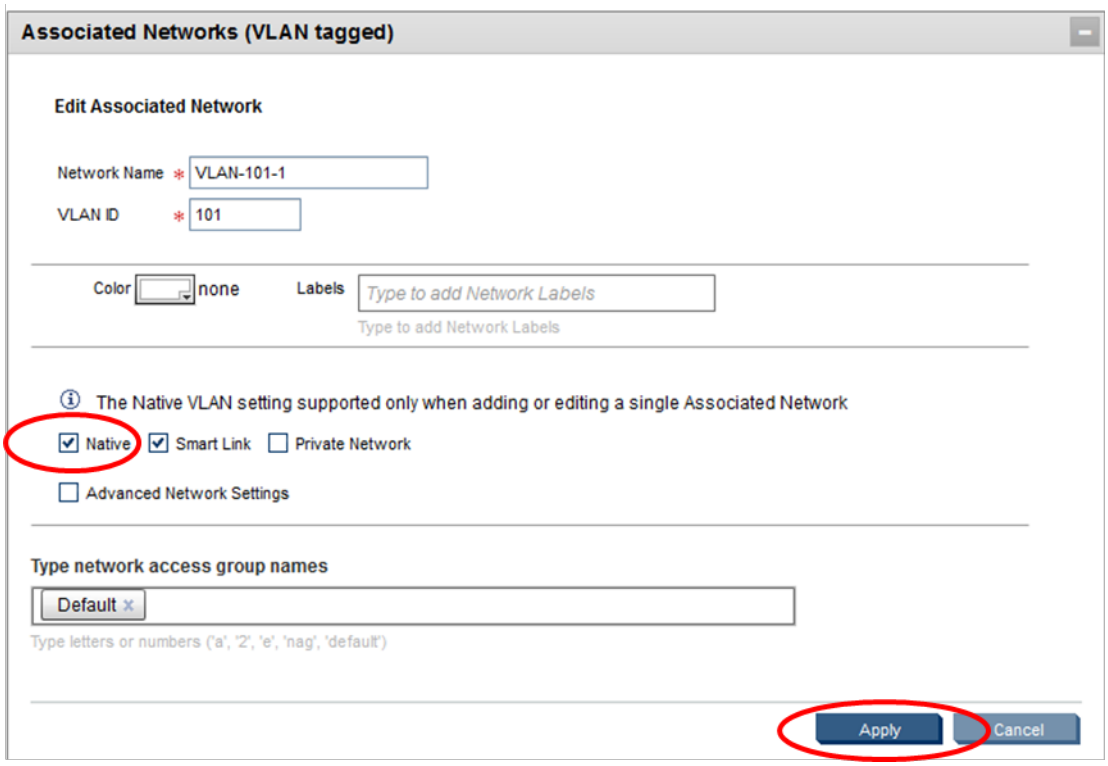

Figure 13 SUS with VLAN-101-1 as Native/Default VLAN

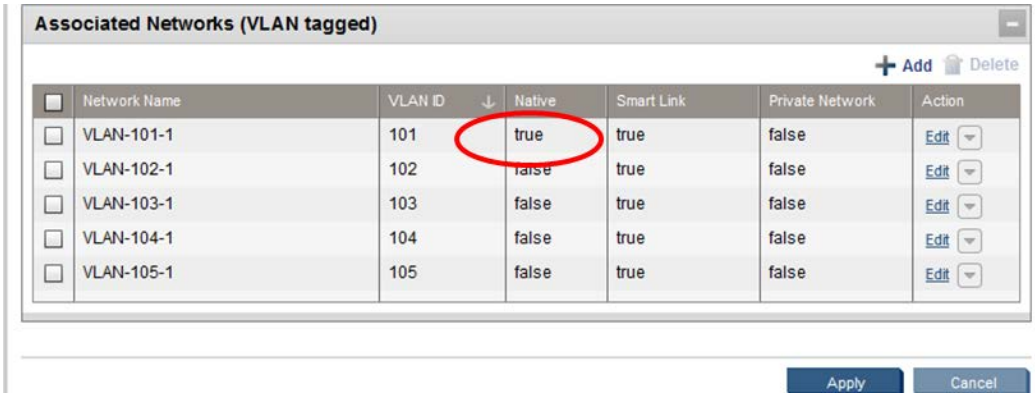

5) Define a new Shared Uplink Set (VLAN-Trunk-2) (Copying a Shared Uplink Set). The second Shared Uplink Set could be created in the same manner as VLAN-Trunk-1 however; VC now provides the ability to COPY a VC Network or Shared Uplink Set

```
In the VC GUI screen, select “Shared Uplink Sets” in the left
pane, in the right pane VLAN-Trunk-1 will be displayed, left
click “VLAN-Trunk-1”, it will appear as blue, right click and
select “COPY”
Edit the Settings as shown below
The new SUS name will be VLAN-Trunk-2
Replace the last instance of 1 with a suffix of 2
In step 3, ADD uplinks X5 and X6 from Bay 2
Click “OK’
```

Figure 14 Adding VLAN-Trunk-2 SUS using copy

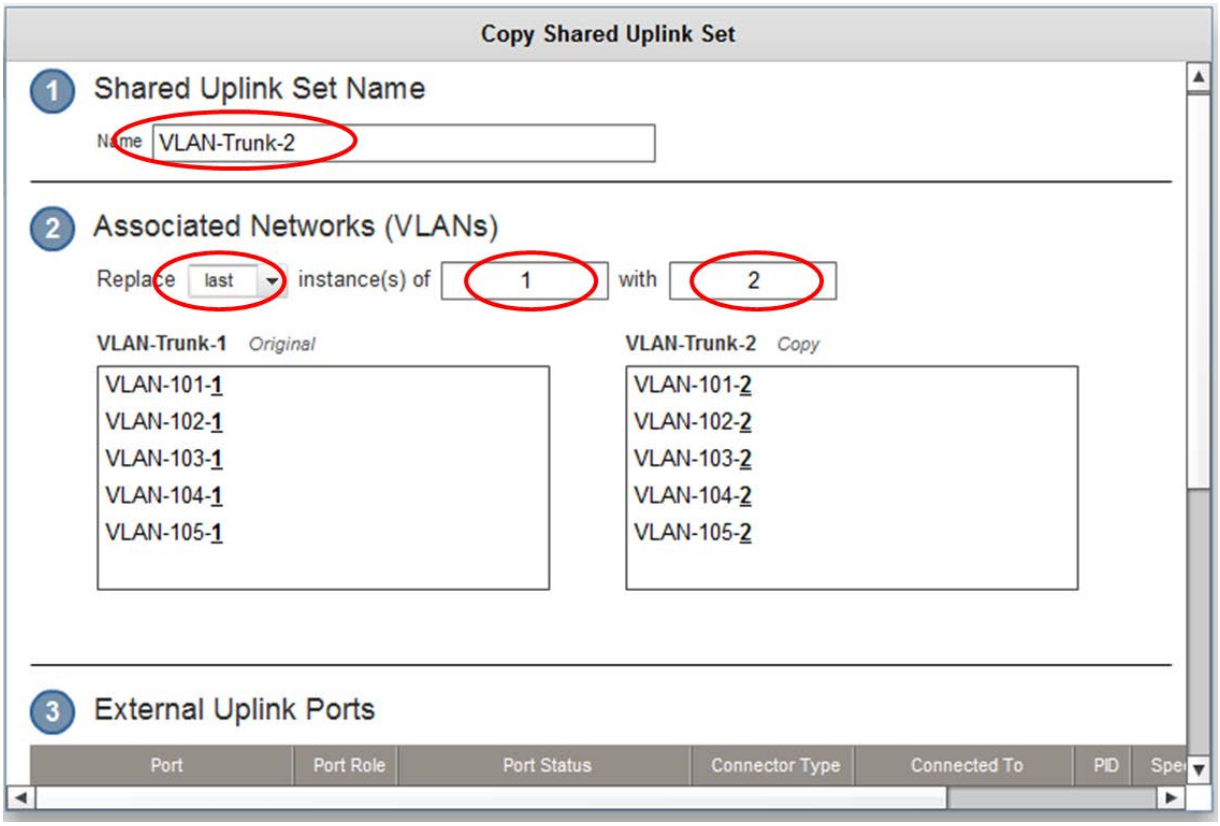

Note: In this scenario we have created two independent Share Uplink Sets (SUS), each originating from the opposite FlexFabric Modules, by doing so we provide the ability to create separate and redundant connections out of the Virtual Connect domain. When we create the server profiles, you will see how the NICs will connect to VLANs accessed through the opposite VC module, which provides the ability to create an Active / Active uplink scenario. Alternatively, we could have created a single SUS and assigned both sets of these uplink ports to the same SUS, however, this would have provided an Active/Standby uplink scenario.

6) Define a vMotion network. This network will provide the framework for allowing east to west vMotion traffic within the enclosure itself.

In the VC GUI screen, select “Ethernet Networks” under “Connections” in the left pane.
Click the plus symbol to add a new network.

Figure 15 Adding vMotion network

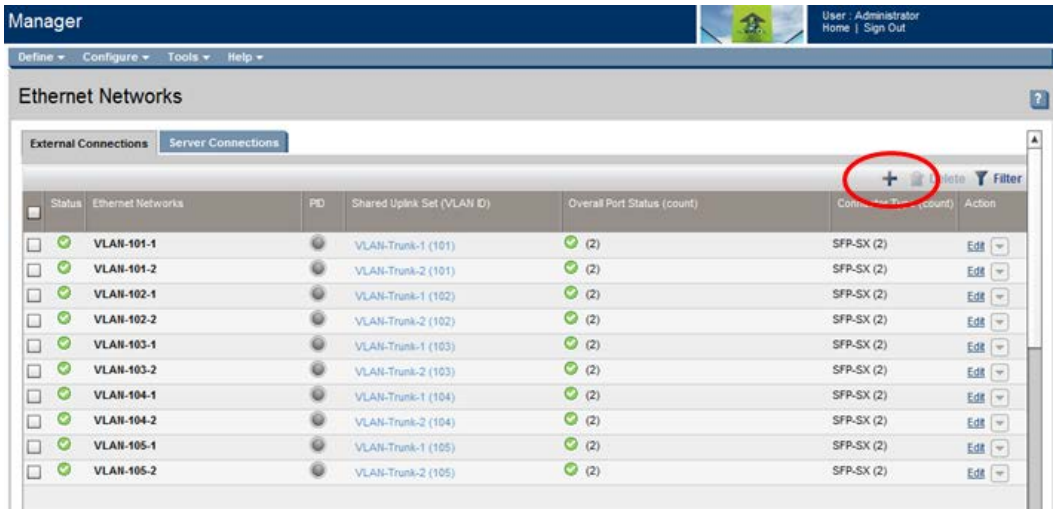

Assign it a network name of “vMotion”, and click “Apply”.

Figure 16 Naming vMotion network

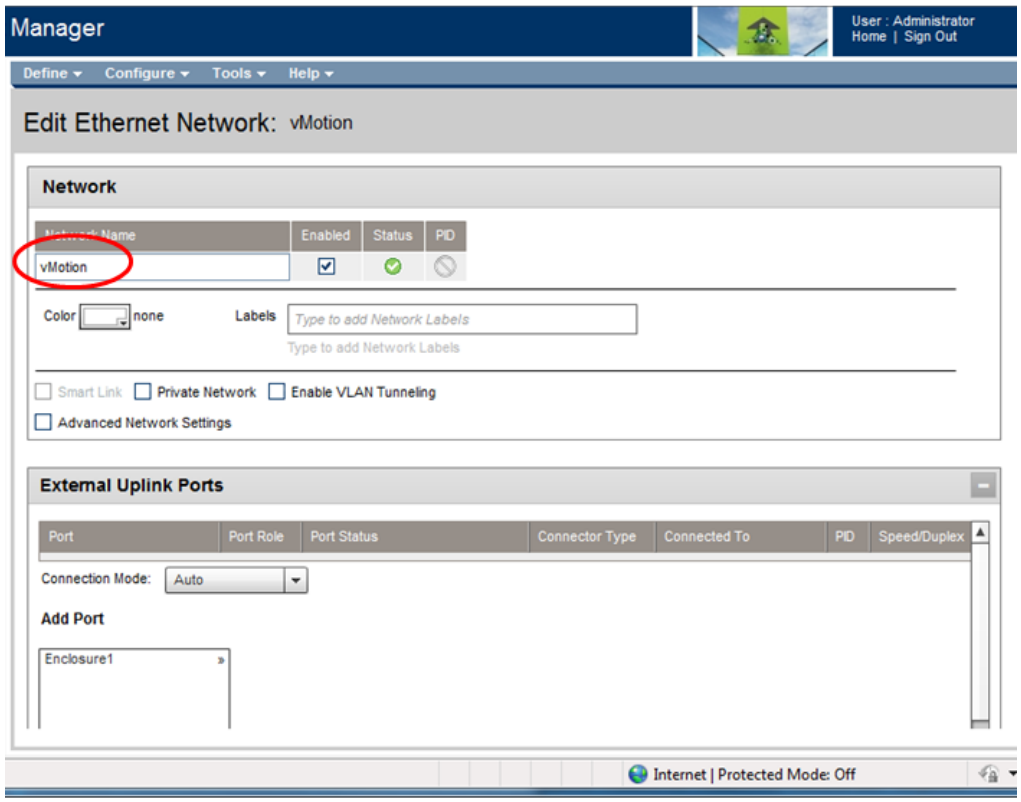

7) Define a new FC SAN Fabric

In the VC GUI screen, select “SAN Fabrics” in the left pane.
Select “Add” in the right pane.
Enter a Fabric name as “SAN-A”
Add ports X1 and X2 from Bay 1 to uplink ports.
In the Fabric Type drop down choose “DirectAttach”
Under Configured Speed choose the appropriate speed. In this example we will choose 8Gb.
Click “Apply”

Figure 17 Creating SAN-A fabric

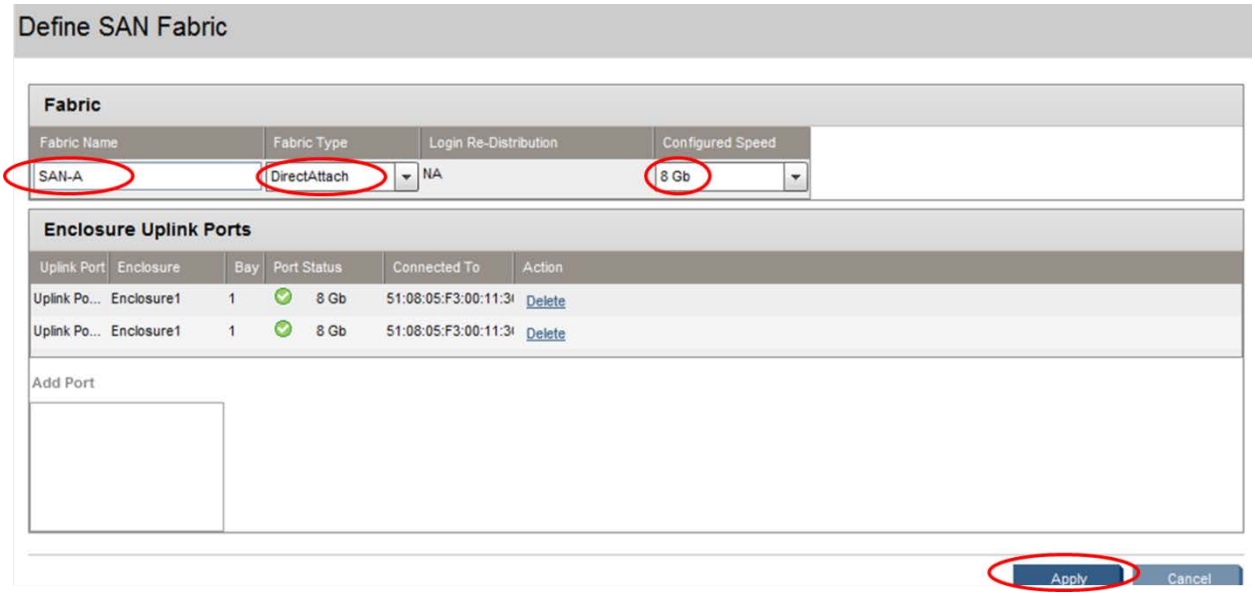

8) Add the second SAN Fabric by again clicking “Add”

Enter a Fabric name as “SAN-B”
Add ports X1 and X2 from Bay 2 to uplink ports.
In the Fabric Type drop down choose “DirectAttach”
Under Configured Speed choose the appropriate speed. In this example we will choose 8Gb.
Click “Apply”

Figure 18 Created SAN fabrics

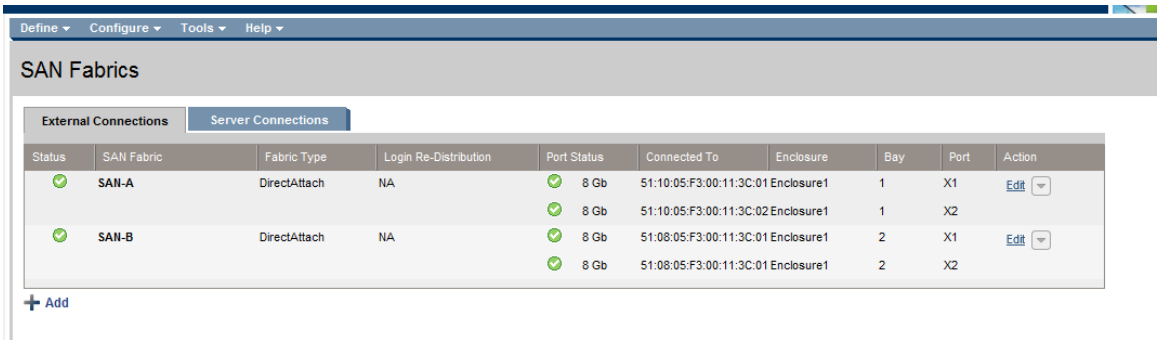

9) Define a Server Profile. In this example we will create a server profile with two server NICs

In the VC GUI screen, select “Define”, then “Server Profile”
Create a server profile called “Profile-1”
In the Ethernet Adapter Connections click the plus symbol twice to add 2 more networks.

Figure 19 Adding 2 more networks to server Profile-1

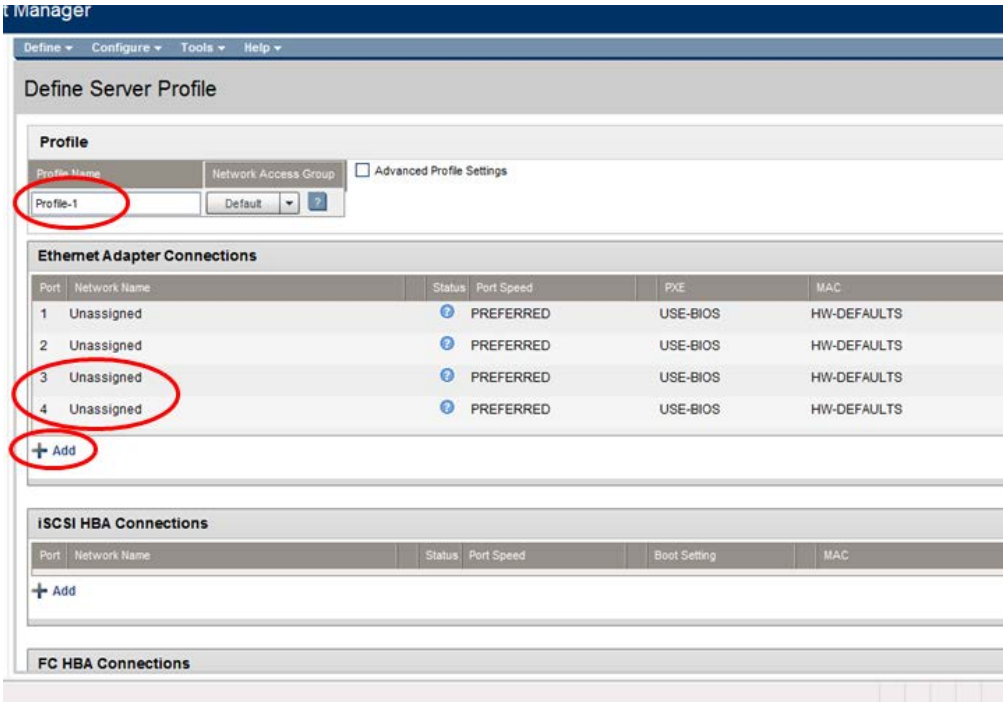

In the Ethernet Adapter Connections section click “unassigned” for Port 1, and then choose “multiple networks” from the drop down box
Select all VLANs with suffix of -1 and drag and drop to window pane to the right
Select “untagged” for VLAN-101-1
Click “OK”

Figure 20 adding Ethernet connections to server profile App-1

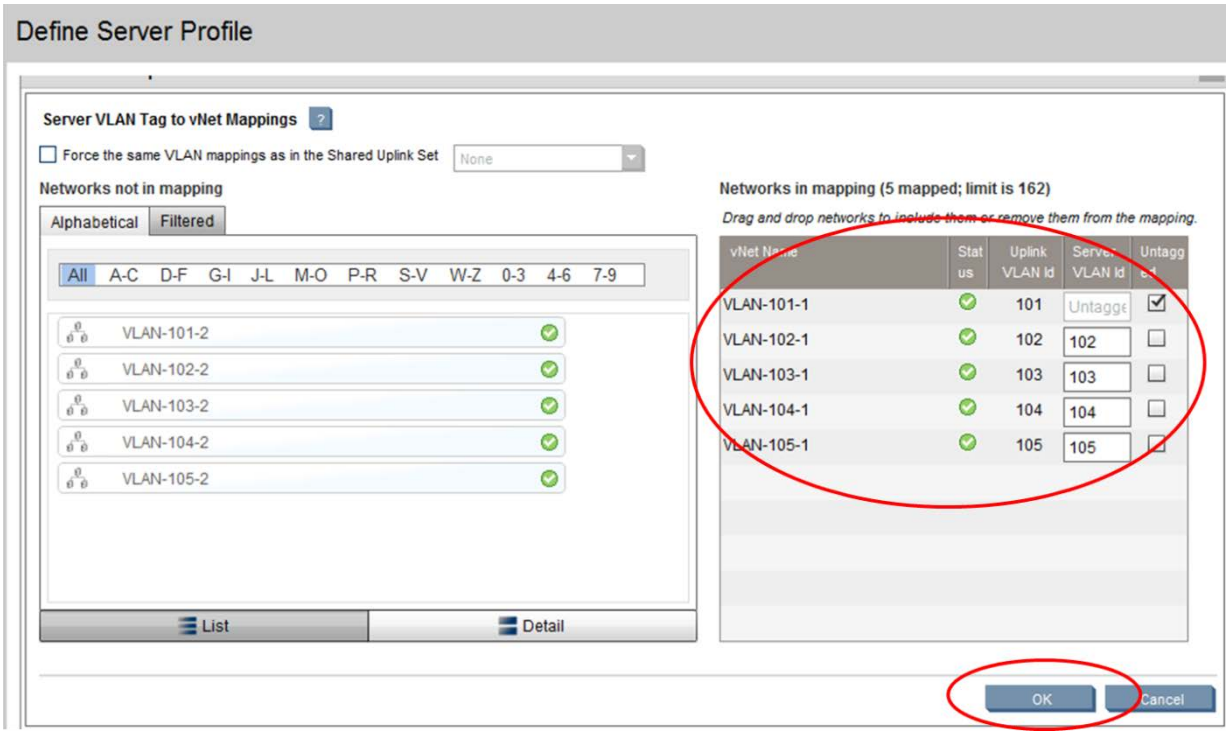

In the Ethernet Adapter Connections section click “unassigned” for Port 2, and then choose “multiple networks” from the drop down box
Select all VLANs with suffix of -2 and drag and drop to window pane to the right
Select “untagged” for VLAN-101-2
Click “OK”
Click “Preferred” under Port Speed for both ports and choose “Custom”
Configure custom port speed to 4Gb
Click “OK”

Figure 21 Custom port speed for Ethernet ports 1 and 2

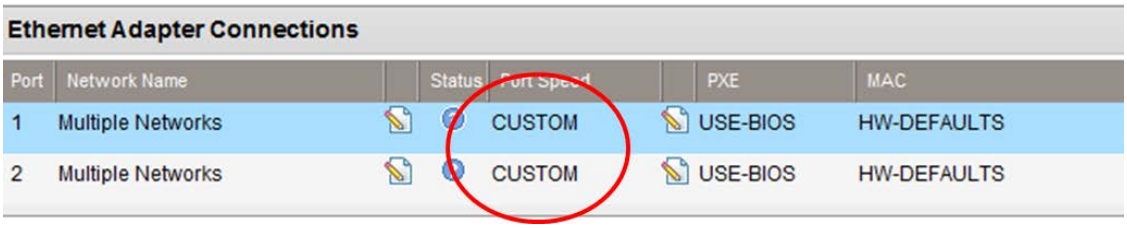

In the Ethernet Adapter Connections section click “unassigned” for Port 3, and then choose “select a network” from the drop down box
Select “vMotion” and click “ok”
In the Ethernet Adapter Connections section click “unassigned” for Port 4, and then choose “select a network” from the drop down box
Select “vMotion” and click “ok”

Click “Preferred” under Port Speed for both ports 3 and 4 and choose “Custom”
Configure custom port speed on ports 3 and 4 to 2Gb
Click “OK”

Figure 22 Ethernet setting on server profile

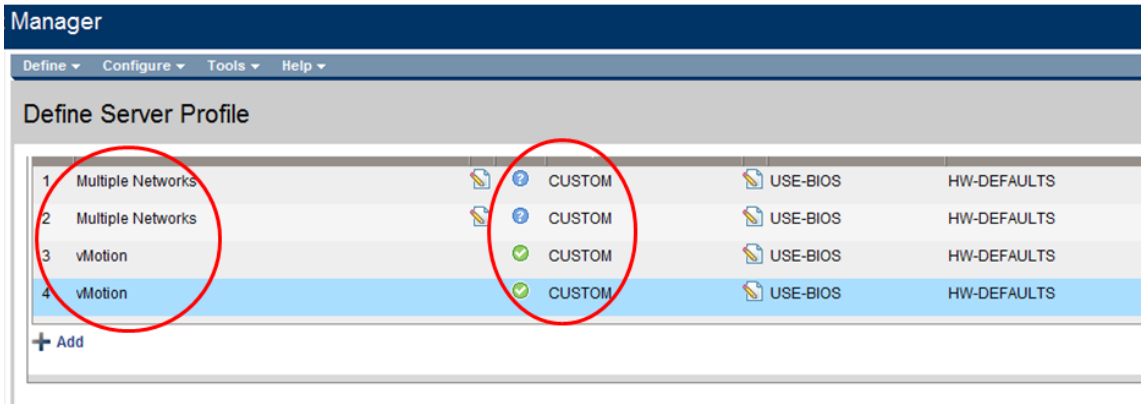

Under the FCoE HBA Connections section, select “unassigned” and select “SAN-A” for Bay 1, and “SAN-B” for Bay 2

Figure 23 Assigning SAN-A and SAN-B to Bay 1

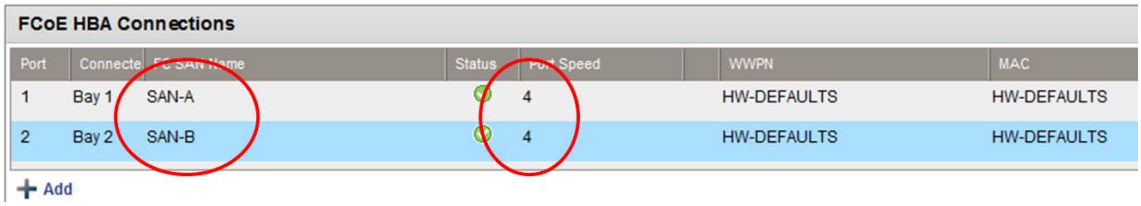

Do not configure FC SAN or iSCSI Connection
In the Assign Profile to Server Bay box, select the server in Bay 1 and then click “apply”

3PAR SAN Deployment Procedures

The following 3PAR deployment procedures listed below are high level baseline configurations intended to provide basic connectivity between the VC and installed 3PAR devices.

- Configure the fabric facing ports
- Create a host
- Define the host persona
- Define the host PWWN
- Create a Common Provisioning Group (CPG)
- Define the CPG disk and RAID type
- Verify the created CPG
- Create a Virtual Volume from the CPG
- Configure the Virtual Volume
- Export the Virtual Volume to the Host

Notes: Complete instructions on setting up 3PAR enclosures is out of scope for this document. The configurations provide an introduction to 3PAR and are basic configurations designed to provide general connectivity. Refer to 3PAR documentation for more info. http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?lang=en&cc=us&taskId=135&prodClassId=-1&contentType=SupportManual&docIndexId=64180&prodTypeId=18964&prodSeriesId=5044394#2

3PAR Port Persistence is not supported when using a Virtual Connect Direct Attach to 3PAR solution. In the event of a node failure within the 3PAR 7200, path failover will occur, however depending on the actual configuration, there will be a I/O pause of roughly 10 seconds.

3PAR SAN Configuration Procedures

The below example steps are screen shots using the 3PAR management console. CLI configurations are supported but are out of scope of this document.

See the 3PAR Management Consoles Use guide for more details:

http://bizsupport1.austin.hp.com/bc/docs/support/SupportManual/c03606437/c03606437.pdf

Figure 24 3PAR management console

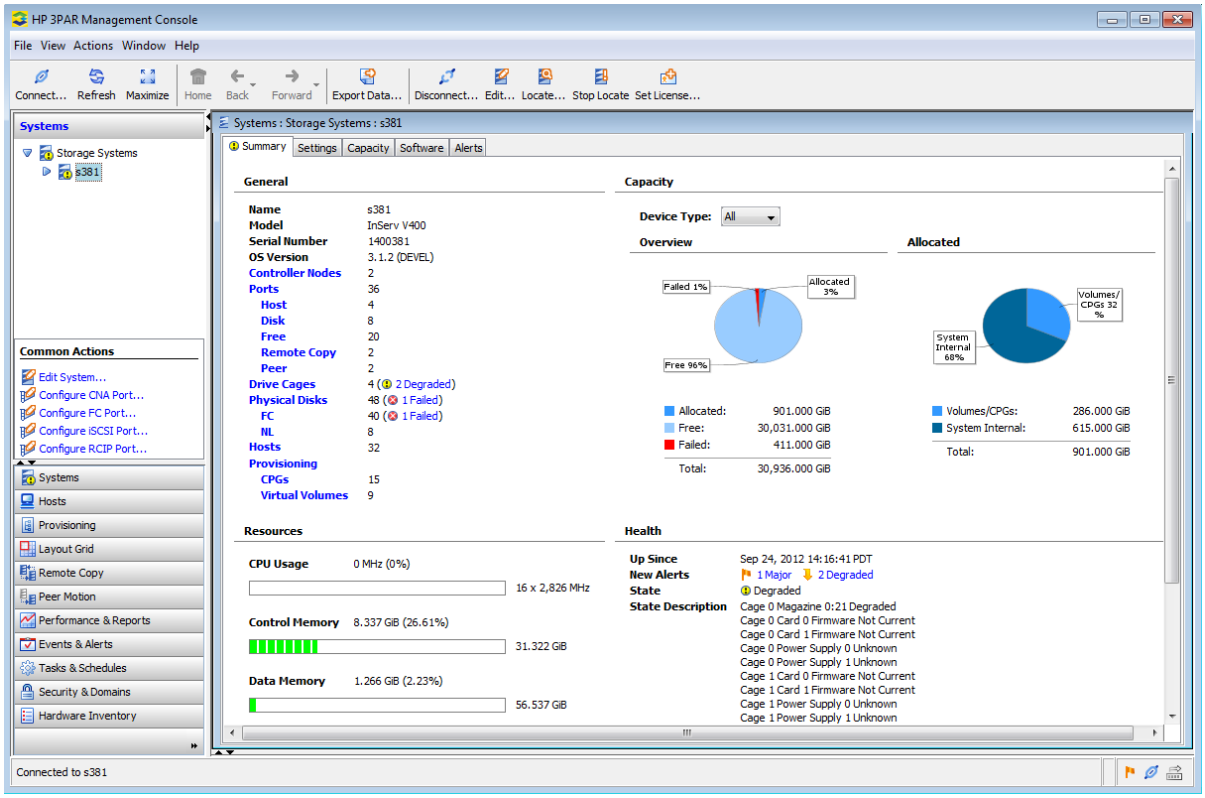

1) Configure the fabric facing ports

In the Manager Pane, click Systems.
In the Management Tree, select the system on which you wish to configure the port.
In the main screen right click the port and choose “Configure”

Figure 25 3PAR configure FC ports

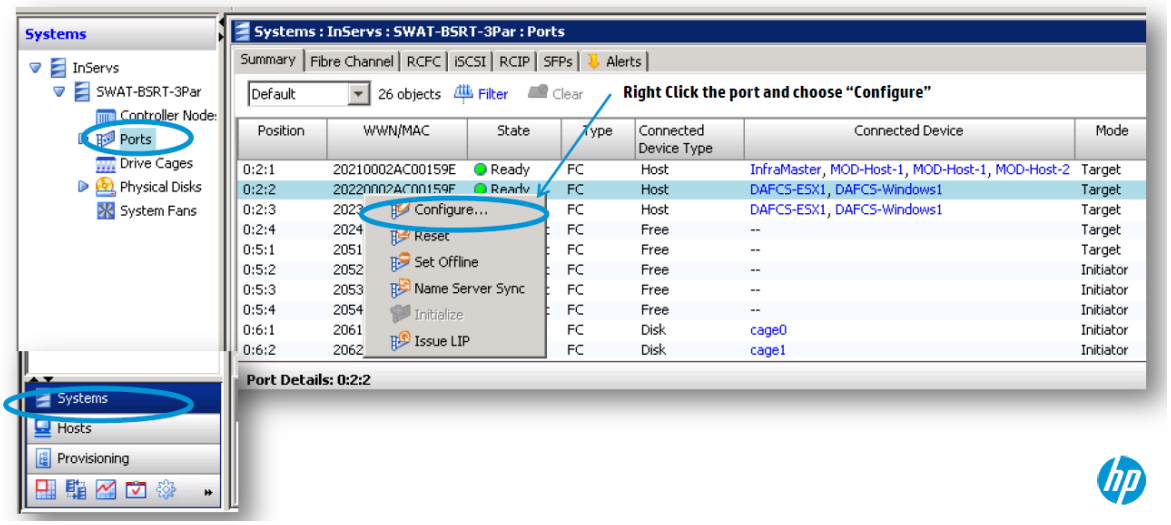

In the Settings group box change the Connection Mode to “Host”
In the Settings group box change the Connection Type to “Point”

Figure 26 3PAR configure FC ports - continued

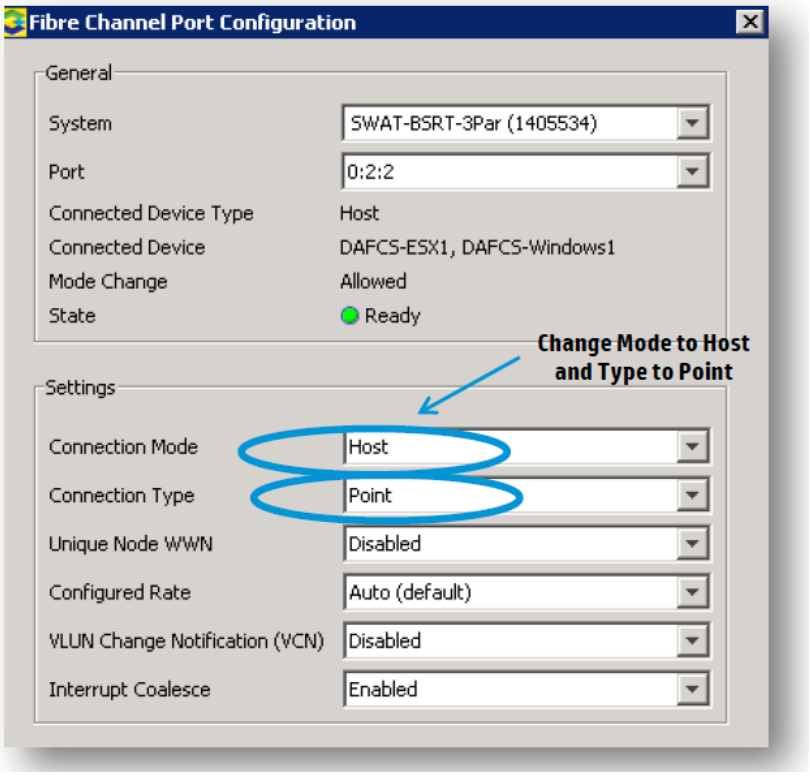

Repeat the same steps for the other 3PAR ports connected to the FlexFabric modules

2) Create a host

Click Hosts in the Manager Pane.
Click Create Host in the Common Actions Panel.

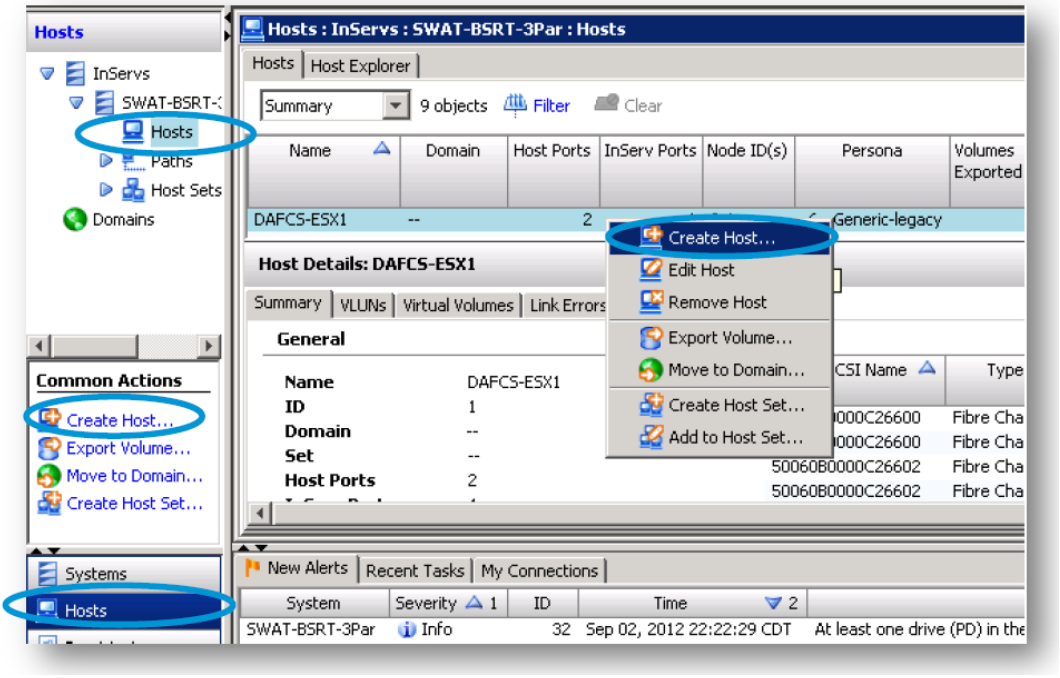

3) Define the host persona

In the General group box: define the “System”, “Domain”, “Name”, “Set Name”, and “Persona”
Click next

Figure 27 3PAR Defining the host persona

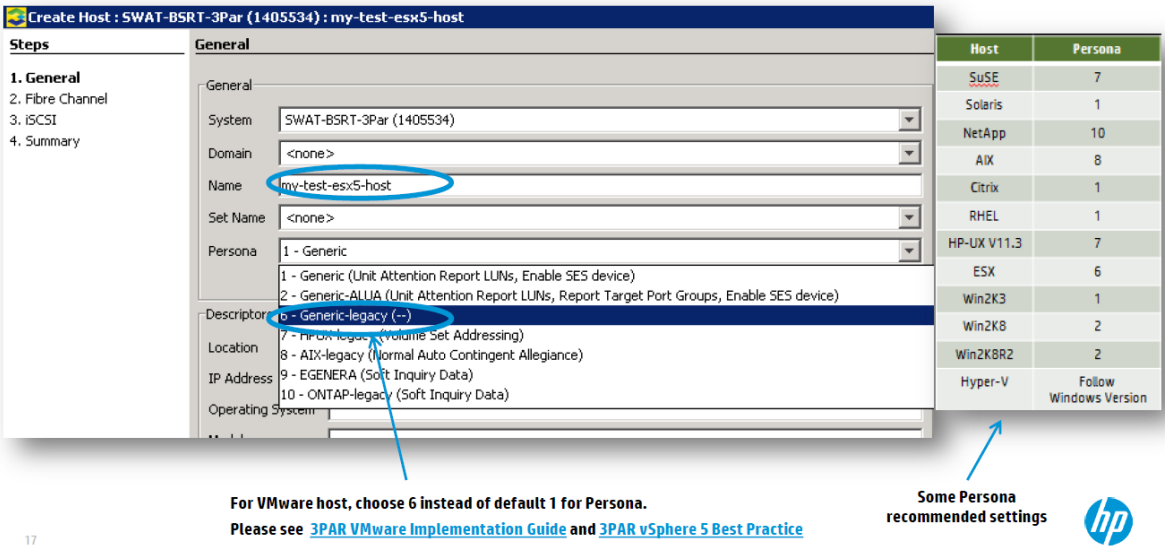

4) Define the host PWWN

To assign available WWNs select one or more WWNs from the Available WWNs list. This list displays WWNs for all physically connected host paths not already assigned to hosts.
Click the left arrow to add the selected WWN(s) to the Assigned WWNs list.
To assign new WWNs, like predefined WWNs that Virtual Connect has assigned, enter the WWN(s) for the host in the New WWN text box and click Assign.
Click next

Figure 28 3PAR assigning WWNs

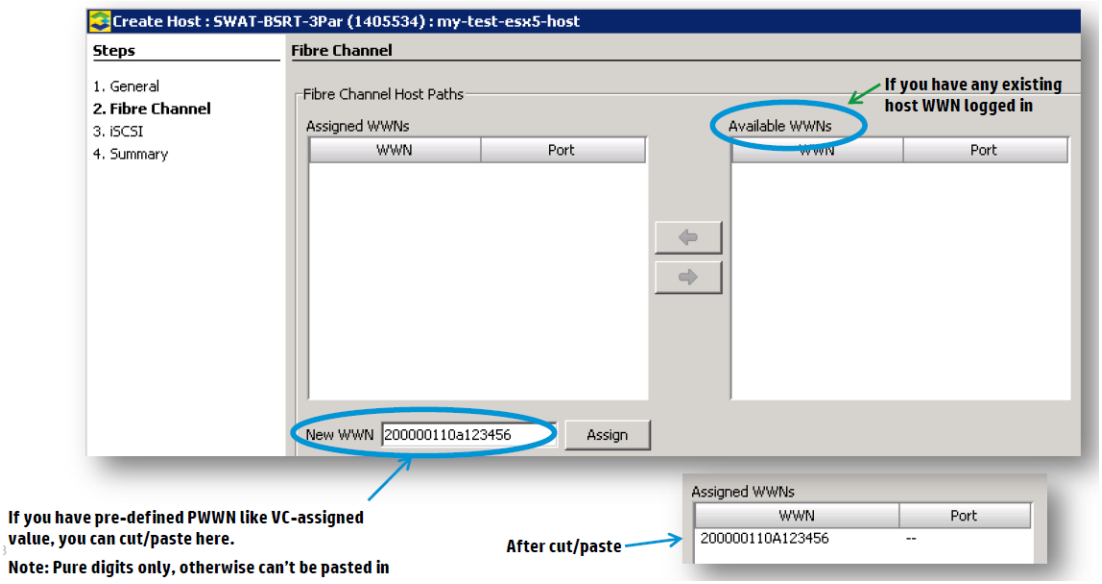

5) Create a Common Provisioning Group (CPG)

```
In the Manager Pane, click Provisioning.
In the Common Actions panel, click Create CPG.
```

Figure 29 3PAR creating CPG

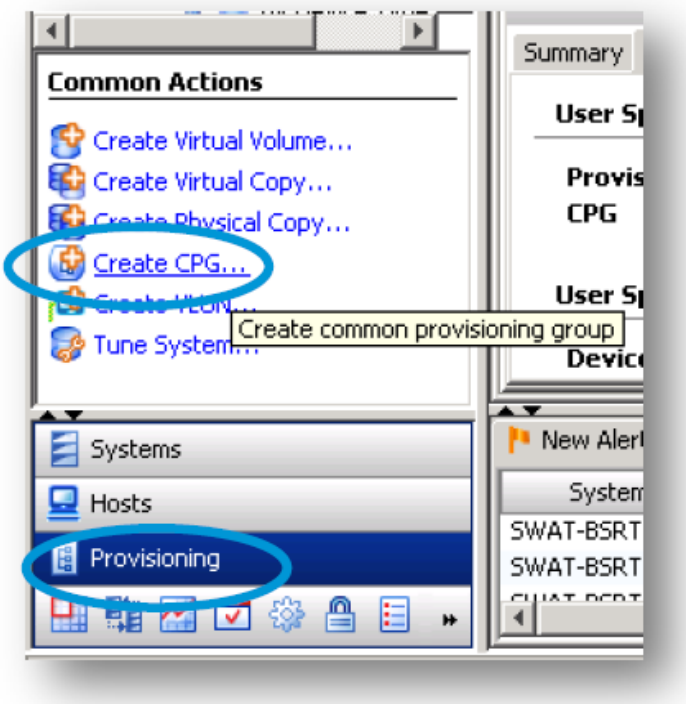

6) Define the CPG disk and RAID type

In the General group box: define the “System”, “Domain”, “Name”, “User Template”, “Device Type”, “Device RPM”, “RAID type”, and “Set Size”
Click next

Figure 30 3PAR Defining the CPG disk and RAID type

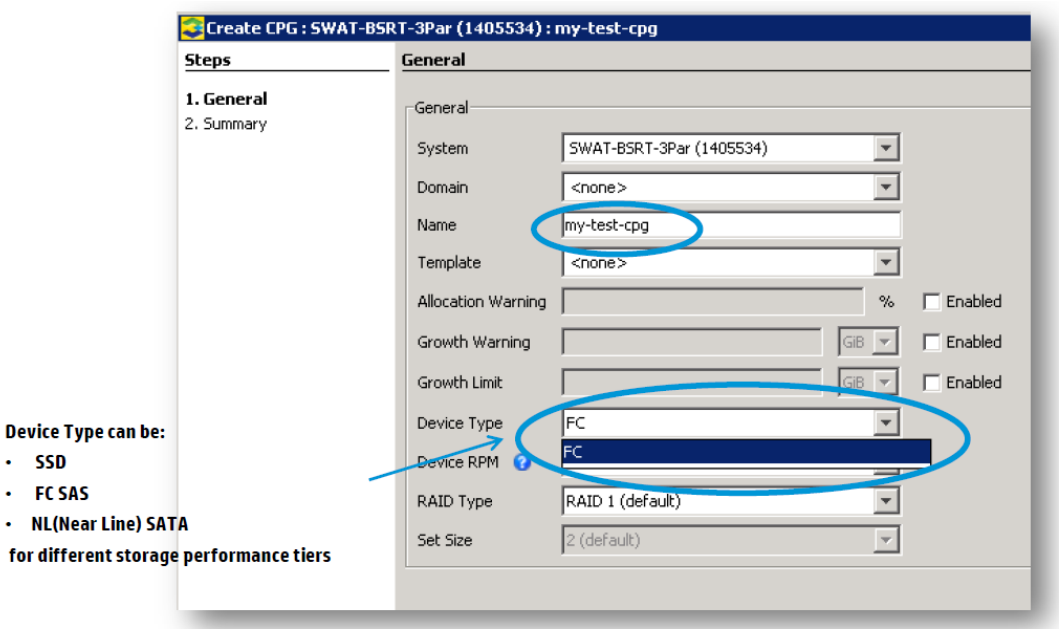

Figure 31 3PAR Defining the CPG disk and RAID type - continued

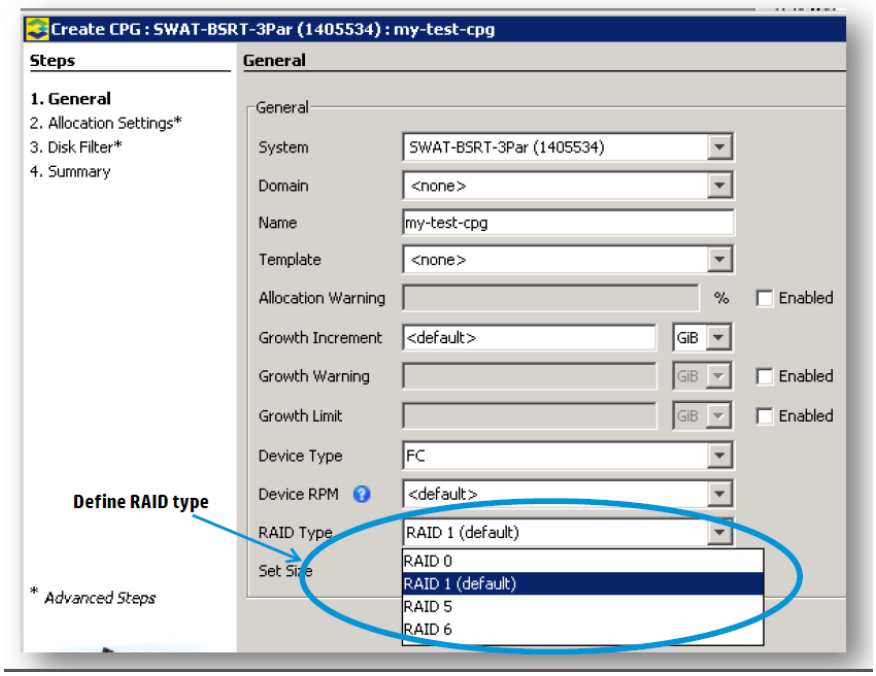

7) Verify the created CPG

Access the CPGs screen through the Provisioning Manager, by clicking the system in the Management Tree containing the CPGs you wish to view.
In the Management Window, click the CPGs tab.
The CPGs tab presents information in a list pane and a detail pane.

Figure 32 3PAR Verifying the created CPG

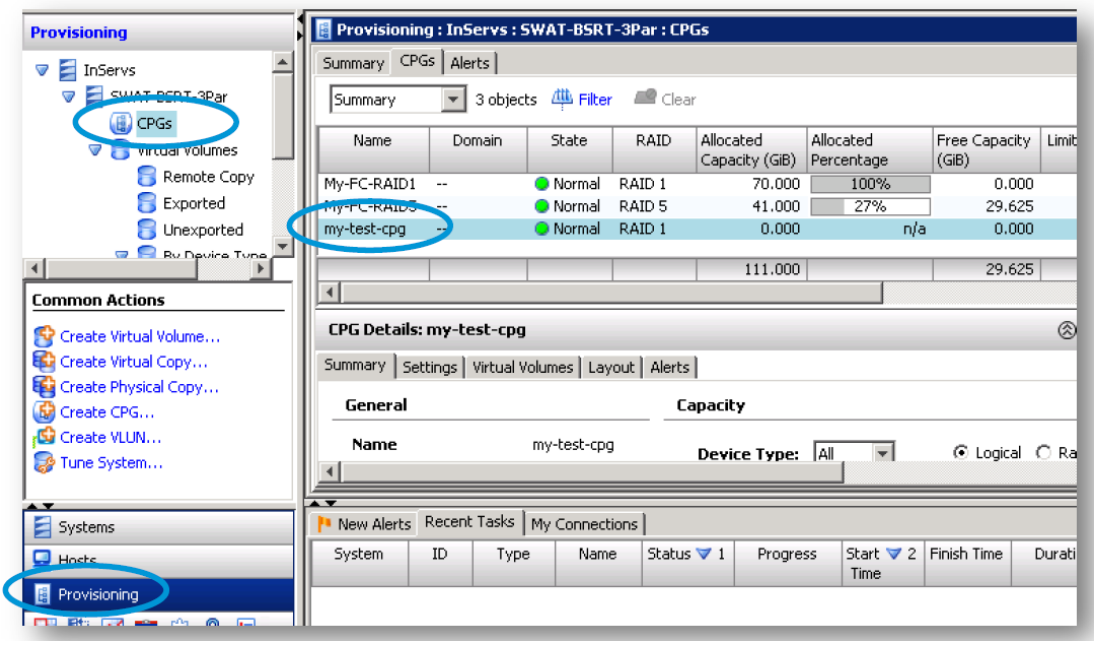

8) Create a Virtual Volume from the CPG

In the Manager Pane, click Provisioning.
Right click the CPG that was just created and choose “Create Virtual Volume”.

Figure 33 3PAR creating a Virtual Volume

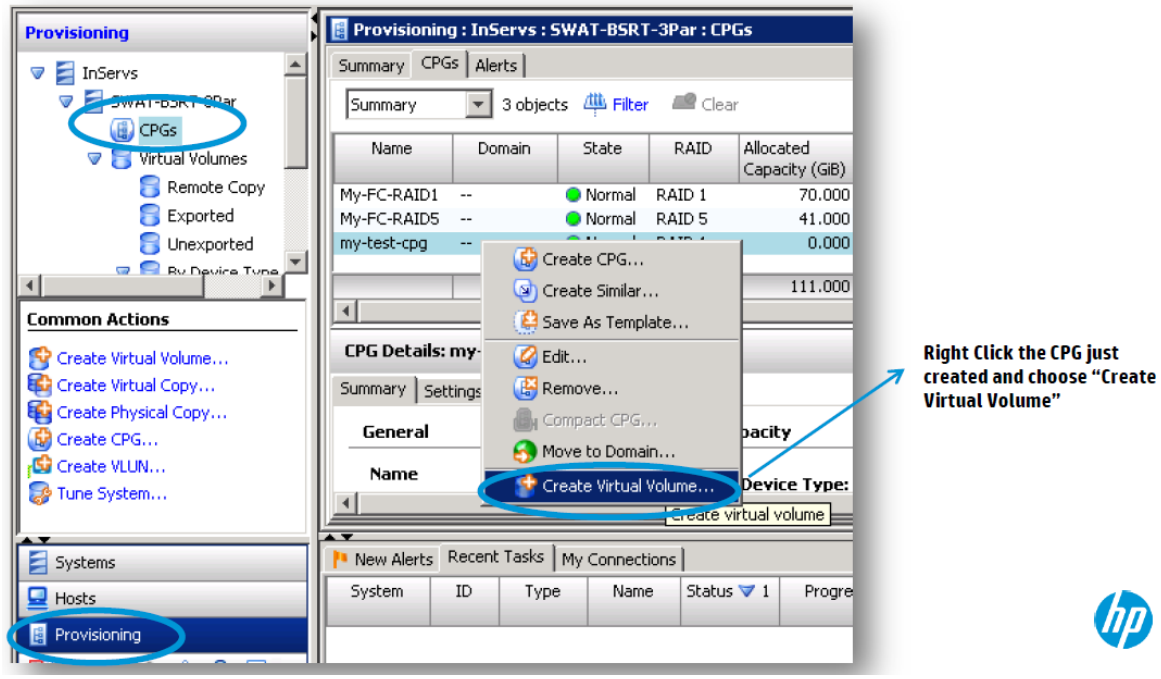

9) Configure the Virtual Volume

In the General group box: define the "System", "Domain", "Name", "ID", "Template", and "Comments"
In the User Space box: define the "Size", "Provisioning", "CPG"
Click next

Figure 34 3PAR Configuring the Virtual Volume

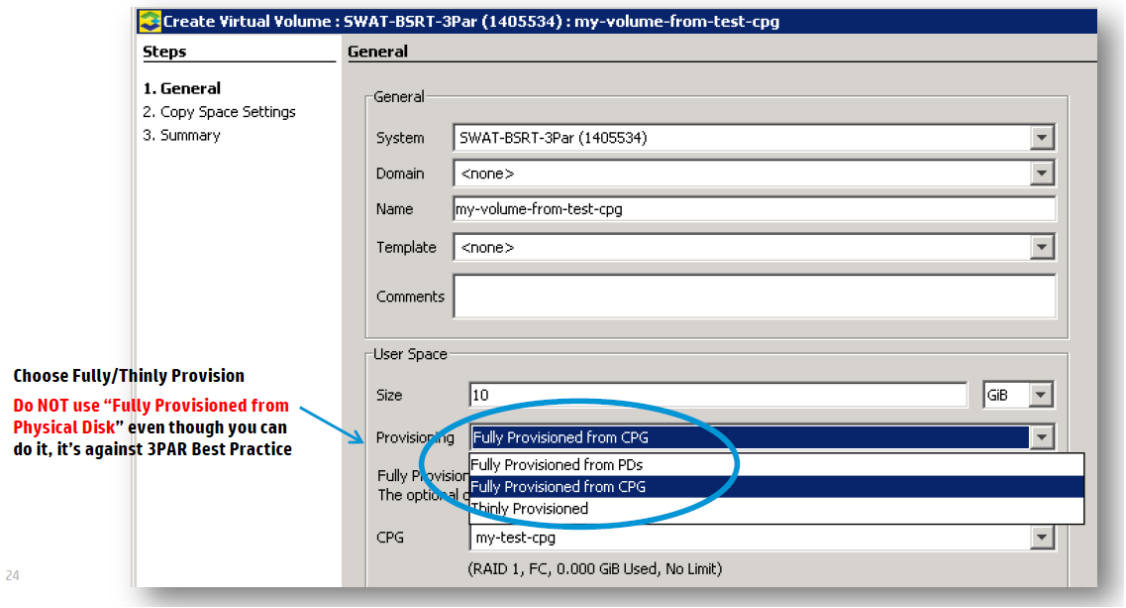

Figure 35 3PAR Configuring the Virtual Volume - continued

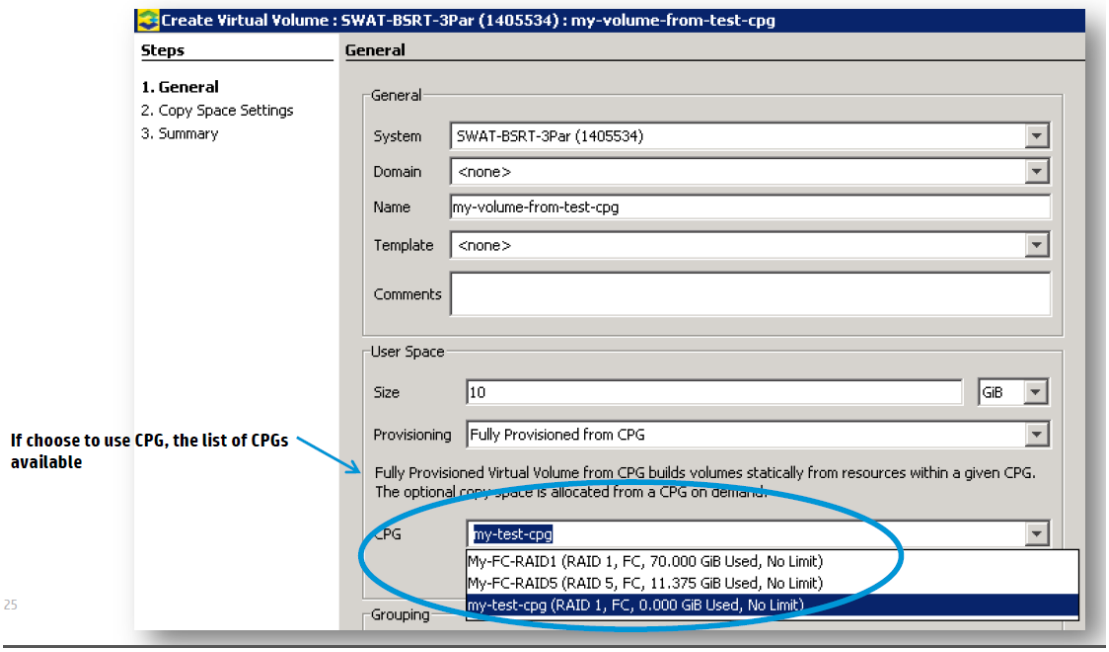

10) Export the Virtual Volume to the Host

Click Hosts or Provisioning in the Manager Pane.
Access the Virtual Volumes tab in the Management Window.
Right-click the volume you wish to export.
In the menu that appears, click Export.

Figure 36 3PAR Exporting the Virtual Volume

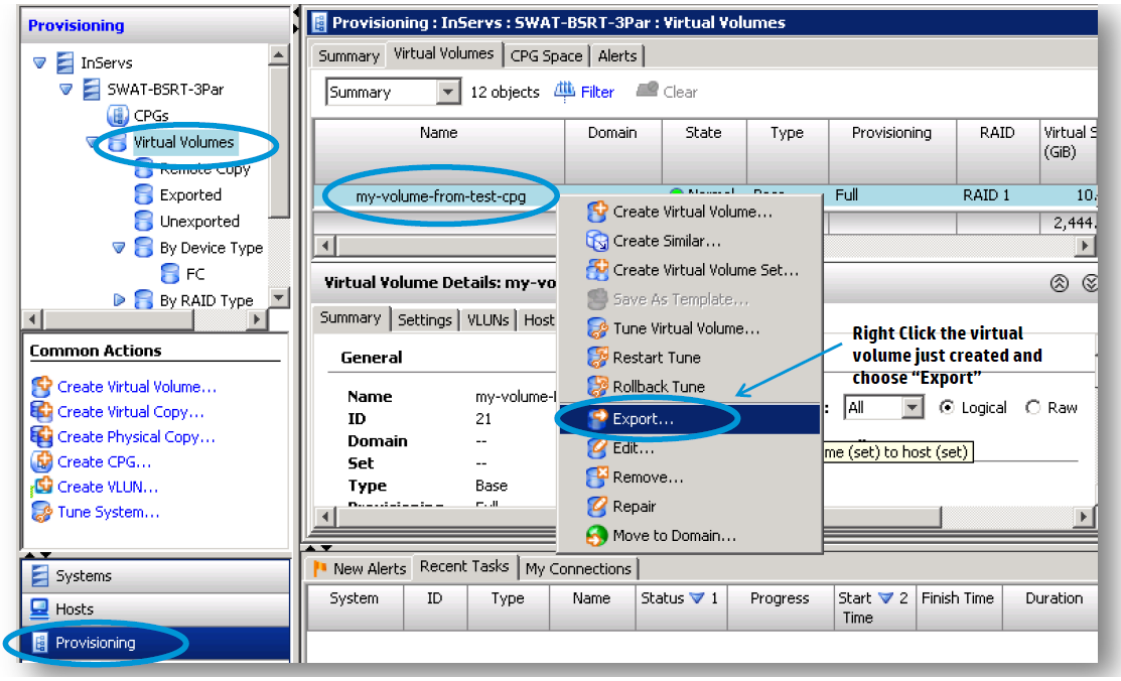

In the General group box, select the “System” and the “Domain”. If not applicable, select <none>.
In the Virtual Volume group box, select either Virtual Volume or Virtual Volume Set. From the virtual volume or virtual volume set list, select the volume(s) to export.

Figure 37 3PAR Exporting the Virtual Volume - continued

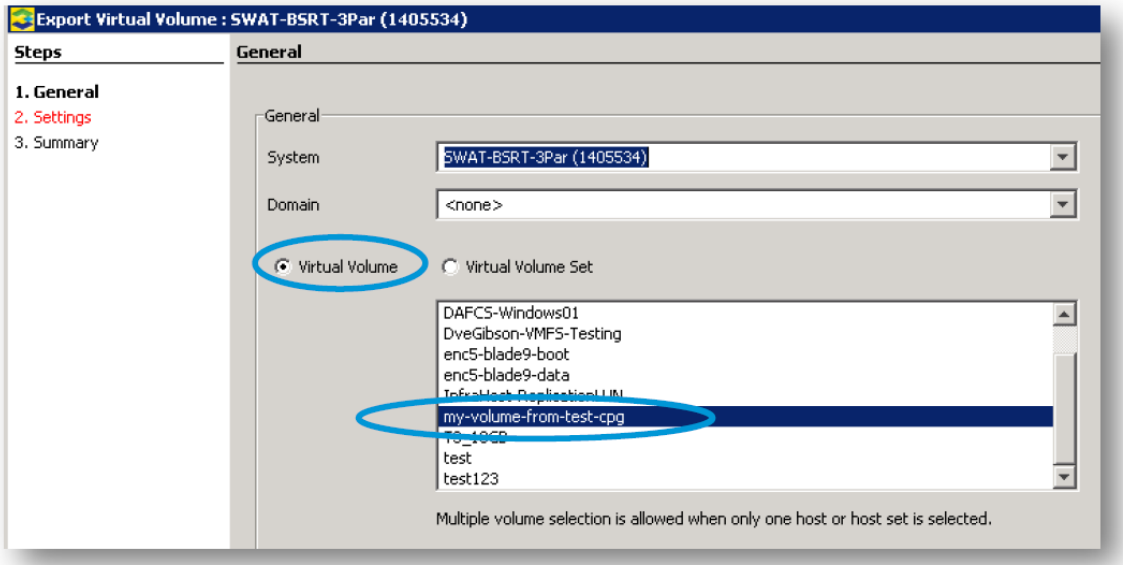

In the Settings group box select either Host or Host Set. From the host or host set list, select the host(s) you wish to export the virtual volumes to.
Optionally choose which local FC port the volume can be exported to, and choose the LUN number.
Click Next to go to the Summary page.

Figure 38 3PAR Exporting the Virtual Volume - continued

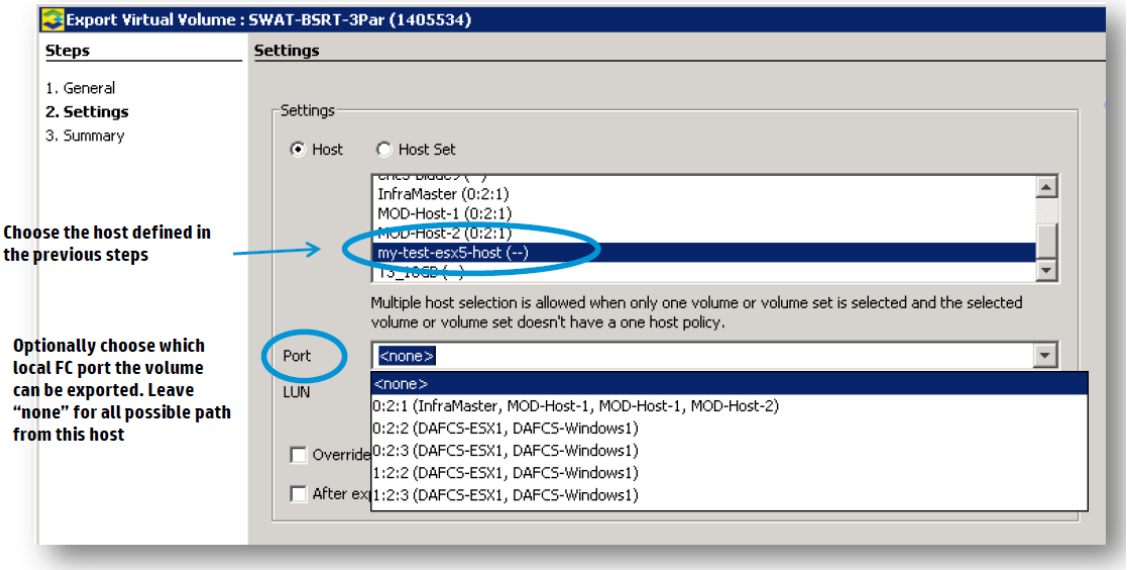

Figure 39 3PAR Exporting the Virtual Volume - continued

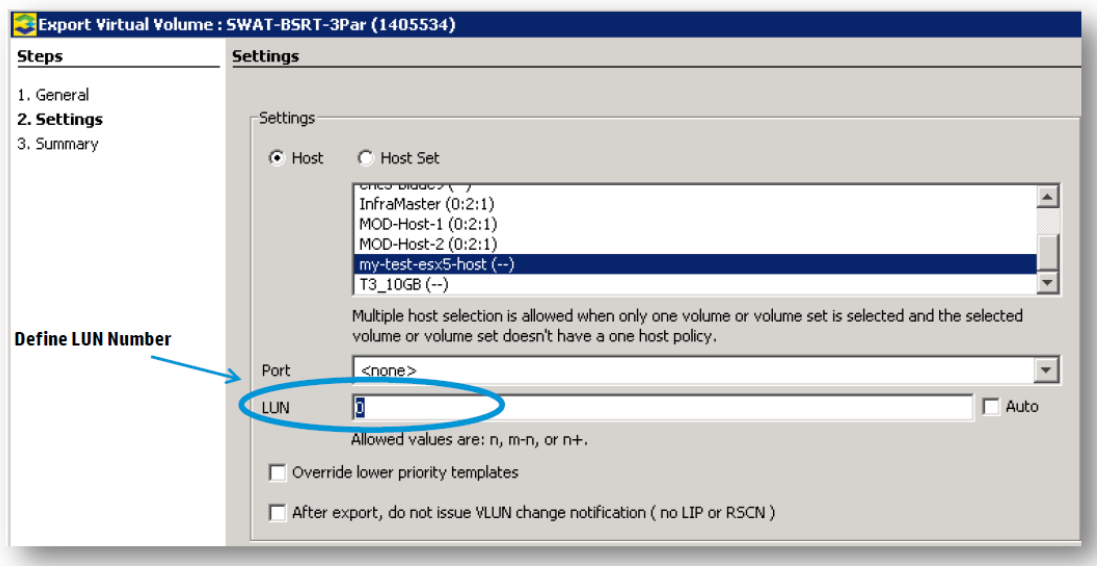

Figure 40 3PAR Exporting the Virtual Volume - continued

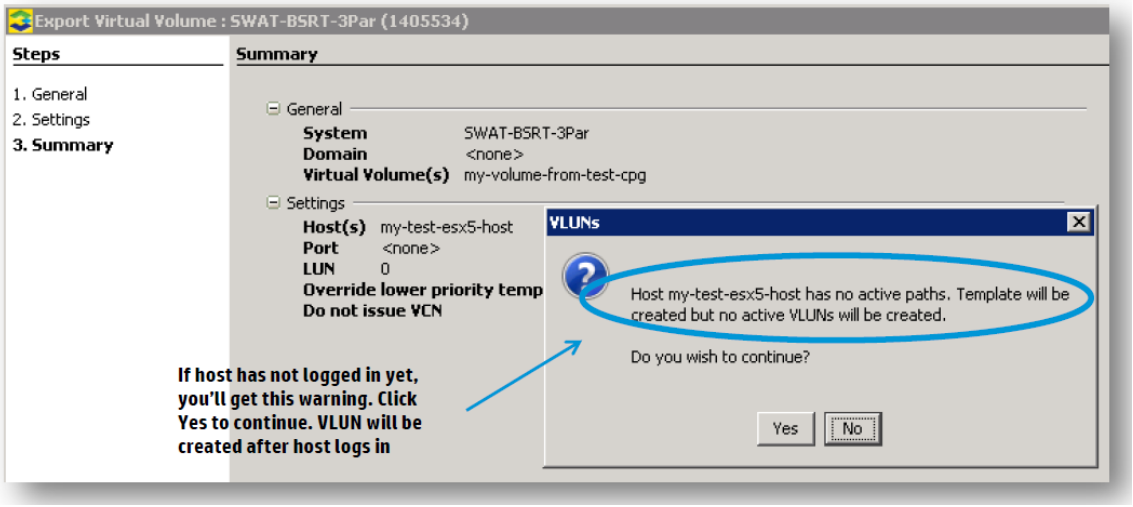

2-Tier Layer 2 ToR Design

The 2-Tier Layer 2 ToR design utilizes a pair of HP 12518 switches at the core which connect to a pair of HP 5900 ToR switches, which then in turn connect to an HP G8 Rack Server.

In this design, the pair of 12518s will still utilize a 40GbE IRF link on the same interfaces as described previously, and act as the gateway for the servers connected to the ToR switches.

The pair of 12518 switches will connect to the pair of 5900 ToR switches using a 120GbE LAG group.

The pair of HP 5900 ToR switches will utilize an 80GbE IRF link between each other and will connect to an HP G8 Rack server using a 20GbE LAG (1*10GbE from each switch). Of course, customers can choose not to implement IRF at the TOR layer if not needed.

From a physical view:

- The 12518 switches will utilize 4 10GbE links for IRF (ports used listed below)
- The pair of 5900s will utilize 2 forty-gig interfaces for the IRF links (ports used listed below)
- The 12518 switches will utilize a 120GbE LAG group (6*10GbE to each 5900) and connect to the pair of 5900s (ports used listed below)
- The pair of 5900s will utilize a 20GbE LACP aggregation group which connects to LOM on an HP G8 Rack Server (ports used listed below)

The scalability of this design is limited by the following aspects:

- As of this writing the 12518 can support a maximum of 288 line rate 10GbE ports or 576 10GbE ports at a 4:1 oversubscription ratio
- Using LEF modules, the 12518 supports 256K MAC address and 64K ARP entries as well as 1M IPv4 FIB entries
- The 12500 series of switches can support 4 chassis IRF for a total of 1,152 line rate 10GbE ports or 2,304 10GbE ports at a 4:1 oversubscription ratio
- The maximum number of Link Aggregation Groups (LAGs) for the 12500 switch is 240 with 12*n ports per LAG (n=number of 12500 chassis in IRF)
- The 5900 supports 128K MAC address entries
- The 5900 supports 16K ARP entries
- Each of the 5900 40GbE interfaces can also support four 10GbE connections
- The maximum number of Link Aggregation Groups (LAGs) for the 5900 is 128 with 16 ports per LAG

Taking in the above criteria, an IRF pair of 12518s with line rate 10GbE ports can accommodate:

- 2,256 physical servers at a 4:1 over-subscription ratio
    - 288 line rate 10GbE ports per chassis *2=576
    - Minus 4 ports per chassis for IRF=568
    - Minus 2 port per chassis for WAN links=564
    - 12 uplinks to each 5900 TOR - 564/12=47
    - 47 switches *48 ports of 10GbE=2,256
        - This solution assumes the 4 40GbE ports have been converted to 16 10GbE ports. 12 of those ports are used of uplinks, leaving 4 10GbE ports which can be used for IRF to another 5900, if needed. If the 40GbE ports are not converted, then this solution supports 1,692 physical servers at 3:1 over-subscription ratio (36*47=1,692)
- If the 5900 TOR switches aren't configured for IRF, the option is to either dual home each server which will reduce the number to 1,128 physical servers or 2,256 single homed servers

Four 12518 with line rate 10GbE ports can accommodate

- 4,464 physical servers at 4:1 oversubscription
    - 288 line rate 10GbE ports per chassis *4=1,152
    - Minus 32 ports for IRF=1,120
    - Minus 1 port per chassis for WAN links=1,116
    - 12 uplinks to each 5900 TOR 1116/12=93
    - 93 switches *48 ports of 10GbE=4,464
        - This solution assumes the 4 40GbE ports have been converted to 16 10GbE ports. 12 of those ports are used of uplinks, leaving 4 10GbE ports which can be used for IRF to another 5900, if needed. If the 40GbE ports are not converted, then this solution supports 3,348 physical servers at 3:1 over-subscription ratio (93*36=3,348)

Note: When scaling this design for virtualized environments, please note that the least common denominator is the 5900 TOR switch that supports 128K MAC entries. With 2 chassis IRF each server could support up to 56 VMs while a 4 chassis design would be limited to 28 VMs per server.

Figure 41 2-Tier Layer 2 Design

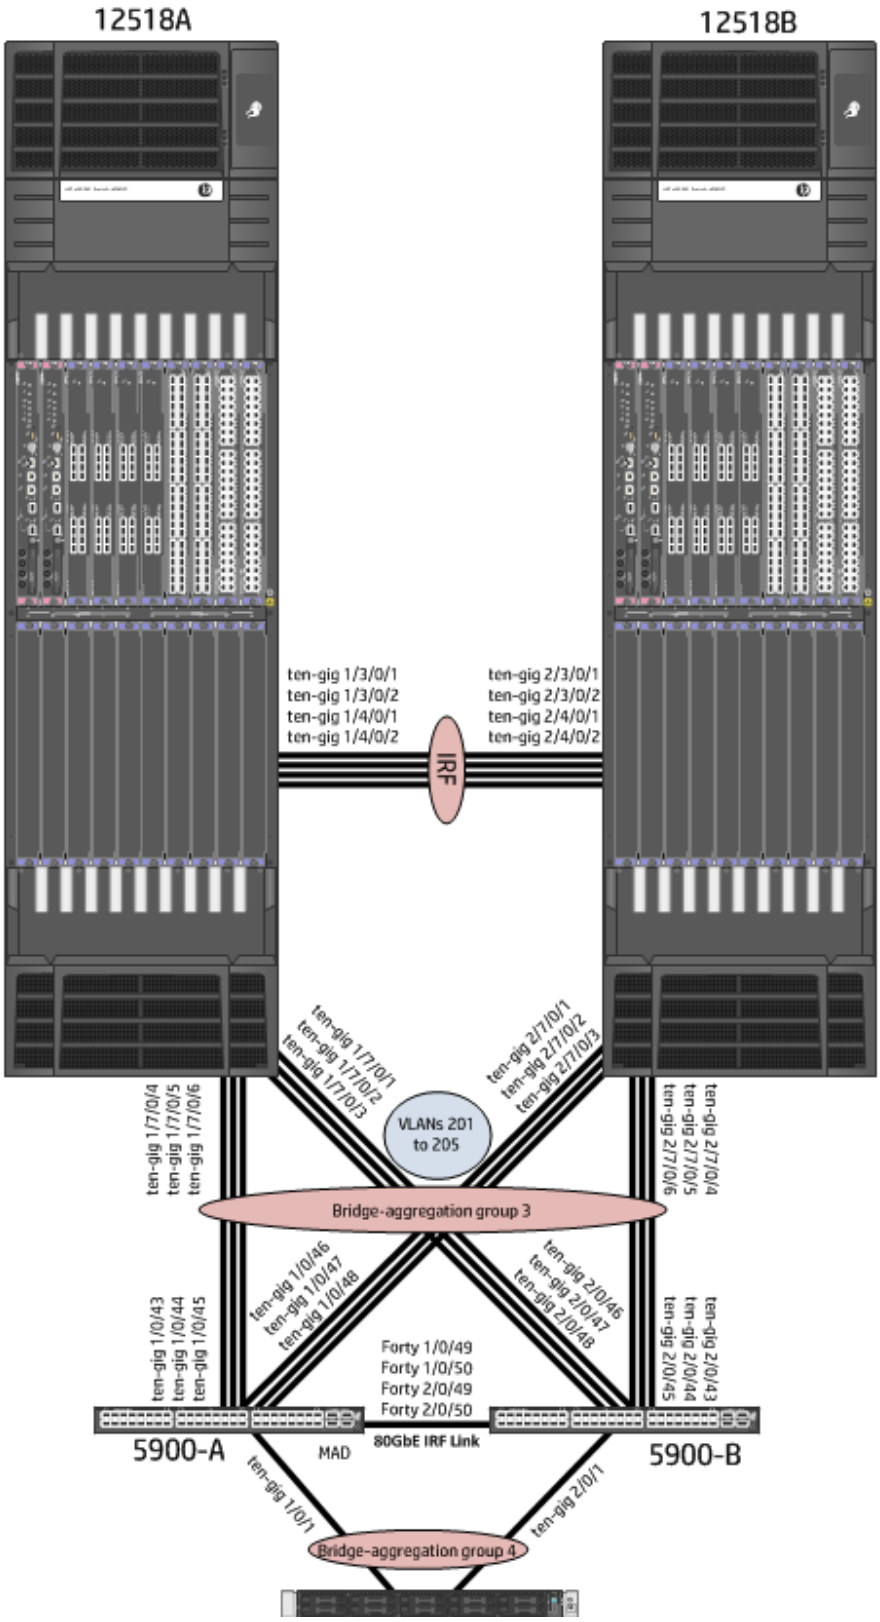

Table 4 2-Tier Layer 2 Design Cabling Information

Link type	Device	Port	Link Type	Remote Device	Port
Bridge Aggregation group 3	12518A	1/7/0/1	10GbE	5900-B	2/0/46
		1/7/0/2			2/0/47
		1/7/0/3			2/0/48
		1/7/0/4		5900-A	1/0/43
		1/7/0/5			1/0/44
		1/7/0/6			1/0/45
	12518B	2/7/0/1	10GbE	5900-A	1/0/46
		2/7/0/2			1/0/47
		2/7/0/3			1/0/48
		2/7/0/4		5900-B	2/0/43
		2/7/0/5			2/0/44
		2/7/0/6			2/0/45
IRF	5900-A	1/0/49	40GbE	5900-B	2/0/49
		1/0/50			2/0/50
Bridge Aggregation group 4	5900-A	1/0/1	10GbE	Rack Server	LOM1
	5900-B	2/0/1	10GbE	Rack Server	LOM2

From a logical view:

- VLANs 201, 202, 203, 204, and 205 will be configured to pass through bridge aggregation groups 3 and 4 to the server as tagged VLANs

Table 5 2-Tier Layer 2 Design VLANs and IP Addresses

Device	VLAN	IP Address	Uplink	Details
12518A (and B)	VLAN201	10.10.201.1 / 24	Bridge-aggregation group 3	Trunk (tagged)
	VLAN202	10.10.202.1 / 24		
	VLAN203	10.10.203.1 / 24		
	VLAN204	10.10.204.1 / 24		
	VLAN205	10.10.205.1 / 24		
5900-A (and B)	VLAN201	N/A	Bridge-aggregation	Trunk (tagged)

	VLAN202 VLAN203 VLAN204 VLAN205		group 3 and 4	
Rack Server	VLAN201 VLAN202 VLAN203 VLAN204 VLAN205	N/A	Bridge-aggregation group 4	LOM1 and 2

LAN Deployment Procedure

This section describes the configuration steps to deploy the HP 12518 and 5900 platforms using IRF for the 2-Tier Layer 2 design. Depending on buffering and oversubscription ratio requirements, different TOR switches can be used.

The LAN procedures listed are in addition to the procedures listed previously in the 1-Tier BladeServer Design section.

- Rack the 5900 switches and install all necessary physical components to establish a functional HP 5900
- Configure IRF on the 5900s using the 40GbE ports
  - Configure IRF
  - Configure LACP MAD detection
- Configure VLANs
  - Create VLAN, names, descriptions
  - Configure VLAN interface IP addresses
  - Undo shutdown VLAN interfaces
  - Configure VLAN MTU, if necessary
- Create and establish uplinks and downlinks
  - Configure Link Aggregation Control Protocol (LACP) for each set of uplinks and downlinks
  - Configure port descriptions to assist in troubleshooting
  - Configure load sharing

LAN Configuration Procedures

Rack the 5900 switches and install all necessary physical components to establish a functional HP 5900

- Detailed descriptions on how to install and rack the HP 5900 switches is out of scope of this document, however a chart of the high level steps that should be followed is listed below. Make sure to fully review the HP 5900 installation manual to ensure proper installation.

- http://h20000.www2.hp.com/bizsupport/TechSupport/CoreRedirect.jsp?redirectReason=DocIndexPDF&prodSeriesId=5221896&targetPage=http%3A%2F%2Fbizsupport2.austin.hp.com%2Fbc%2Fdocs%2Fsupport%2FSupportManual%2Fc03189333%2Fc03189333.pdf

Figure 42 5900 Installation flow chart

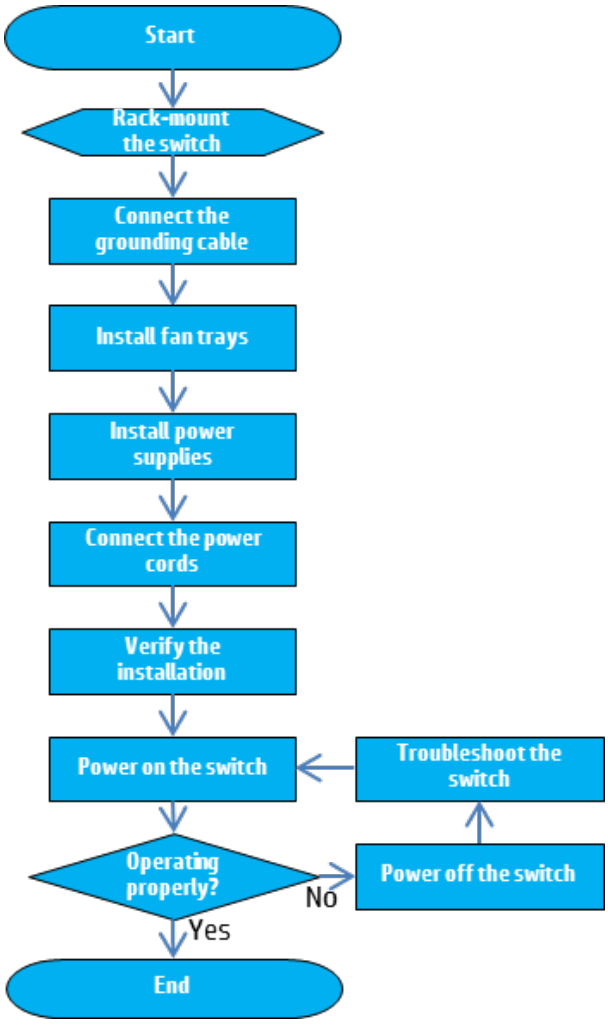

Configure IRF & BFD MAD detection on the 5900s

- The configuration listed refers to the two HP 5900 switches as 5900-A and 5900-B. Once IRF configuration has been established the IRF pair will be referred to as 5900-A
- To offset the risk of IRF virtual device partition, configure MAD to detect multi-active collisions. In this example, LACP MAD is adopted in this example as HP devices are capable of using this method. If other manufacturer TOR switches are used, IRF BFD should be used.

Configure 5900-A

1) Set the member ID of 5900-A to 1, then save the configuration

```
irf member 1
save
```

2) Change the IRF priority to 30, then save the configuration. The higher IRF priority will determine the IRF Master

```
irf priority 30
quit
save
```

3) Shutdown the IRF interfaces

```
interface forty-gigabitethernet 1/0/49
  shutdown
interface forty-gigabitethernet 1/0/50
  shutdown
  quit
```

4) Create IRF port 1/2, and bind the physical IRF ports to it

```
irf-port 1/2
  port group interface forty-gigabitethernet 1/0/49
  port group interface forty-gigabitethernet 1/0/50
              quit
```

5) Undo the Shutdown on the IRF interfaces and save the configuration

```
interface forty-gigabitethernet 1/0/49
  undo shutdown
interface forty-gigabitethernet 1/0/50
  undo shutdown
  quit
save

The current configuration will be written to the device. Are you
sure? [Y/N]:y
Please input the file name(*.cfg)[flash:/startup.cfg]
(To leave the existing filename unchanged, press the enter key):
flash:/startup.cfg exists, overwrite? [Y/N]:y
Validating file. Please wait.....................................
The current configuration is saved to the active main board
successfully.
Configuration is saved to device successfully.
```

Configure 5900-B

1) Set the member ID of 5900-B to 2, then save the configuration

```
irf member 2
save
```

2) Specify the priority. If this switch is going to be an IRF slave, there is no need to modify the default priority of 1

3) Shutdown the IRF interfaces

```
interface forty-gigabitethernet 2/0/49
  shutdown
interface forty-gigabitethernet 2/0/50
  shutdown
quit
```

4) Create IRF port 1/1, and bind the physical IRF ports to it

```
irf-port 1/1
  port group interface forty-gigabitethernet 2/0/49
  port group interface forty-gigabitethernet 2/0/50
  quit
```

5) Undo the Shutdown on the IRF interfaces and save the configuration

```
interface forty-gigabitethernet 2/0/49
  undo shutdown
interface forty-gigabitethernet 2/0/50
  undo shutdown
  quit
save
The current configuration will be written to the device. Are you
sure? [Y/N]:y
Please input the file name(*.cfg)[flash:/startup.cfg]
(To leave the existing filename unchanged, press the enter key):
flash:/startup.cfg exists, overwrite? [Y/N]:y
Validating file. Please wait.....................................
The current configuration is saved to the active main board
successfully.
Configuration is saved to device successfully.
```

6) If the remaining 40GbE interfaces are to be used for 10GbE connectivity, use the following command.. After the configuration, the system prompts you to reboot your device, and you must reboot the device to see the four 10-GE interfaces.

```
interface FourtyGigE1/0/51
  using tengige
```

Connect the two switches with the appropriate cables and bring up the IRF fabric

1) Turn off both switches
2) Connect the IRF links
3) Turn on the first switch and wait until it finishes the boot cycle
4) Turn on the second switch
5) When both switches have completed their boot cycles, verify the configuration:

```
display irf
```

```
Switch Role Priority CPU-Mac Description
*+1 Master 30 b8af-6731-c7ad -----
2 Slave 1 b8af-672c-4308 -----
-------------------------------------------------
  * indicates the device is the master.
  + indicates the device through which the user logs in.
  The Bridge MAC of the IRF is: b8af-6731-c77a
  Auto upgrade          : yes
  Mac persistent        : always
  Domain ID             : 1

display irf configuration
MemberID    NewID       IRF-Port1         IRF-Port2
1     1           disable           forty-Gig1/0/49
forty-Gig1/0/50
2     2           forty -Gig2/0/49  disable
forty -Gig2/0/50
```

Configure LACP MAD

1) Set 5900-A domain ID as 2

```
irf domain 2
```

2) This configuration will be using LAG 3 to connect to Comware ToR switches. Enable LACP MAD on LAG3 on the core switches (steps are shown in the link aggregation configuration section below)

Create VLAN-interface and configure its IP address.

Configure 12518A

1) Create VLAN-interface 201 and configure its IP address as 10.10.201.1/24, and undo shutdown

```
vlan 201
  description To ToR
  quit
interface vlan-interface 201
  ip address 10.10.201.1 24
  undo shutdown
  quit
```

2) Create VLAN-interface 202 and configure its IP address as 10.10.202.1/24, and undo shutdown

```
vlan 202
  description To ToR
  quit
interface vlan-interface 202
  ip address 10.10.202.1 24
  undo shutdown
  quit
```

3) Create VLAN-interface 203 and configure its IP address as 10.10.203.1/24, and undo shutdown

```
vlan 203
```

```
  description To ToR
  quit
interface vlan-interface 203
  ip address 10.10.203.1 24
  undo shutdown
  quit
```

4) Create VLAN-interface 204 and configure its IP address as 10.10.204.1/24, and undo shutdown

```
vlan 204
  description To ToR
  quit
interface vlan-interface 204
  ip address 10.10.204.1 24
  undo shutdown
  quit
```

5) Create VLAN-interface 205 and configure its IP address as 10.10.205.1/24, and undo shutdown

```
vlan 205
  description To ToR
  quit
interface vlan-interface 205
  ip address 10.10.205.1 24
  undo shutdown
  quit
```

Configure 5900A

6) Create VLAN-interface 201 to 205 on 5900-A, and undo shutdown.

```
vlan 201 to 205
interface vlan-interface 201
  undo shutdown
  quit
interface vlan-interface 202
  undo shutdown
  quit
interface vlan-interface 203
  undo shutdown
  quit
interface vlan-interface 204
  undo shutdown
  quit
interface vlan-interface 205
  undo shutdown
  quit
```

Link Aggregation Configuration Procedure

- Configure 1 dynamic Link-Aggregation group (LAG 3) on 12518A and 5900-A and assign the appropriate ports
- Configure 1 dynamic Link-Aggregation groups (LAG 4) on 5900-A and assign appropriate ports

- Configure each LAG as trunk ports and ensure that they each carry tagged VLANs 201 to 205
- Enable LACP MAD on 12518

Configure 12518A

1) Create Layer 2 aggregate interface Bridge-Aggregation 3 on 12518A, configure the link aggregation mode as dynamic, and enable LACP MAD

```
interface bridge-aggregation 3
  description LAG to ToR-A and B
  link-aggregation mode dynamic
  mad enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 1]:
The assigned domain ID is: 1
Info: MAD LACP only enable on dynamic aggregation interface
  quit
```

2) Configure Layer 2 aggregate interface Bridge-Aggregation 3 as a trunk port and assign VLANs 201 to 205

```
interface bridge-aggregation 3
  port link-type trunk
  port trunk permit vlan 201 to 205
Please wait... Done.
  quit
```

3) Assign ports to link aggregation group 3 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 1/7/0/1
  description To 5900-B ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/7/0/2
  description To 5900-B ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/7/0/3
  description To 5900-B ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/7/0/4
  description To 5900-A ToR Rack 01/02
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/7/0/5
  description To 5900-A ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/7/0/6
  description To 5900-A ToR
  port link-aggregation group 3
```

```
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/7/0/1
  description To 5900-A ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/7/0/2
  description To 5900-A ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/7/0/3
  description To 5900-A ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/7/0/4
  description To 5900-B ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/7/0/5
  description To 5900-B ToR
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/7/0/6
  description To 5900-B ToR
  port link-aggregation group 3
  undo shutdown
  quit
```

4) Configure the 12518A to use either the source and destination IP or the source and destination MAC addresses as the global link-aggregation load sharing criteria. In virtual environments, consider using source/destination IP load-sharing as MAC load-sharing may favor a single link. Both options are shown below.

```
link-aggregation load-sharing mode source-ip destination-ip
link-aggregation load-sharing mode source-mac destination-mac
```

Configure 5900A

1) Create Layer 2 aggregate interface Bridge-Aggregation 3 on 5900-A, and configure the link aggregation mode as dynamic

```
interface bridge-aggregation 3
  link-aggregation mode dynamic
  quit
```

2) Configure Layer 2 aggregate interface Bridge-Aggregation 3 as a trunk port and assign it to VLANs 201 to 205

```
interface bridge-aggregation 3
```

```
  port link-type trunk
  port trunk permit vlan 201 to 205
Please wait... Done.
  quit
```

3) Assign ports to link aggregation group 3 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 1/0/43
  description To 12518 Core-A
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/44
  description To 12518 Core-A
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/45
  description To 12518 Core-A
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/46
  description To 12518 Core-B
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/47
  description To 12518 Core-B
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/48
  description To 12518 Core-B
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/0/43
  description To 12518 Core-B
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/0/44
  description To 12518 Core-B
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/0/45
  description To 12518 Core-B
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/0/46
  description To 12518 Core-A
  port link-aggregation group 3
  undo shutdown
```

```
  quit
interface Ten-GigabitEthernet 2/0/47
  description To 12518 Core-A
  port link-aggregation group 3
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/0/48
  description To 12518 Core-A
  port link-aggregation group 3
  undo shutdown
  quit
```

4) Configure the 5900 to use either the source and destination IP or the source and destination MAC addresses as the global link-aggregation load sharing criteria. In virtual environments, consider using source/destination IP load-sharing as MAC load-sharing may favor a single link. Both options are shown below.

```
link-aggregation load-sharing mode source-ip destination-ip
link-aggregation load-sharing mode source-mac destination-mac
```

5) Create Layer 2 aggregate interface Bridge-Aggregation 4 on 5900-A, and configure the link aggregation mode as dynamic

```
interface bridge-aggregation 4
  link-aggregation mode dynamic
  quit
```

6) Configure Layer 2 aggregate interface Bridge-Aggregation 4 as a trunk port and assign it to VLANs 201 to 205

```
interface bridge-aggregation 4
  port link-type trunk
  port trunk permit vlan 201 to 205
Please wait... Done.
  quit
```

7) Assign ports to link aggregation group 4 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 1/0/1
  description To Blade Server
  port link-aggregation group 4
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/0/1
  description To Blade Server
  port link-aggregation group 4
  undo shutdown
  quit
```

Rack Server Deployment Procedure Configuration

Since there are many variables when configuring rack servers for specific customer environment, the following items should be considered. Certain setup aspects will vary based on the host OS and are outside the scope of this document.

- iLO management parameters
- Will Intelligent Provisioning be used
- Power options
- NIC options
    - Active/active LACP group
    - Active/Standby NICs

Please refer to the server manuals for your particular server model. An example of the DL380p Gen8 server is linked below.

http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?lang=en&cc=us&contentType=SupportManual&prodTypeId=15351&prodSeriesId=5177957

Rack Server Configuration Procedure

Since rack servers are typically custom made to order, the configuration procedure will be different based on the server deployment. Below is a checklist of items that should be performed. A detailed configuration is outside the scope of this document.

- Install components like memory, NICs , hard drives
- Rack the server
- Connect the power supply
- Connect a terminal, or keyboard, monitor and mouse
- Configure iLO settingsDeploy the host OS using either Intelligent provisioning or
- Once host OS is setup configure NIC options
- Configure storage options

2-Tier Layer 3 ToR Design

The 2-Tier Layer 3 ToR design utilizes the pair of HP 12518 switches at the core which connect to HP 5900 ToR switches, which then in turn connect to HP G8 Rack Servers.

In this design, the pair of 12518s will still utilize a 40GbE IRF link on the same interfaces as described previously, while the Layer 3 routing has been pushed down to the HP 5900 ToR switch. The 5900 will then connect to an HP G8 Rack server using a 10GbE connection. In this design the 5900 ToR switch will be acting as an individual switch, rather than a pair using IRF as previously shown, however, you could deploy pairs using IRF and the configuration would be similar.

This L3 design is a way to overcome the MAC explosion problem that a layer 2 only design exhibits. It also provides fault isolation to the rack level instead of being propagated across the whole DC fabric. Inter-rack routing is accomplished using OSPF.

From a physical view:

- The 12518 switches will utilize 4 10GbE links for IRF (ports used listed below)
- Each physical 12518 switch will utilize a 60GbE LAG group (6*10GbE) to the 5900 (ports used listed below)
- The 5900 will utilize a 10GbE connection to the HP G8 Rack Server (port used listed below)

Figure 43 2-Tier Layer 3 Design

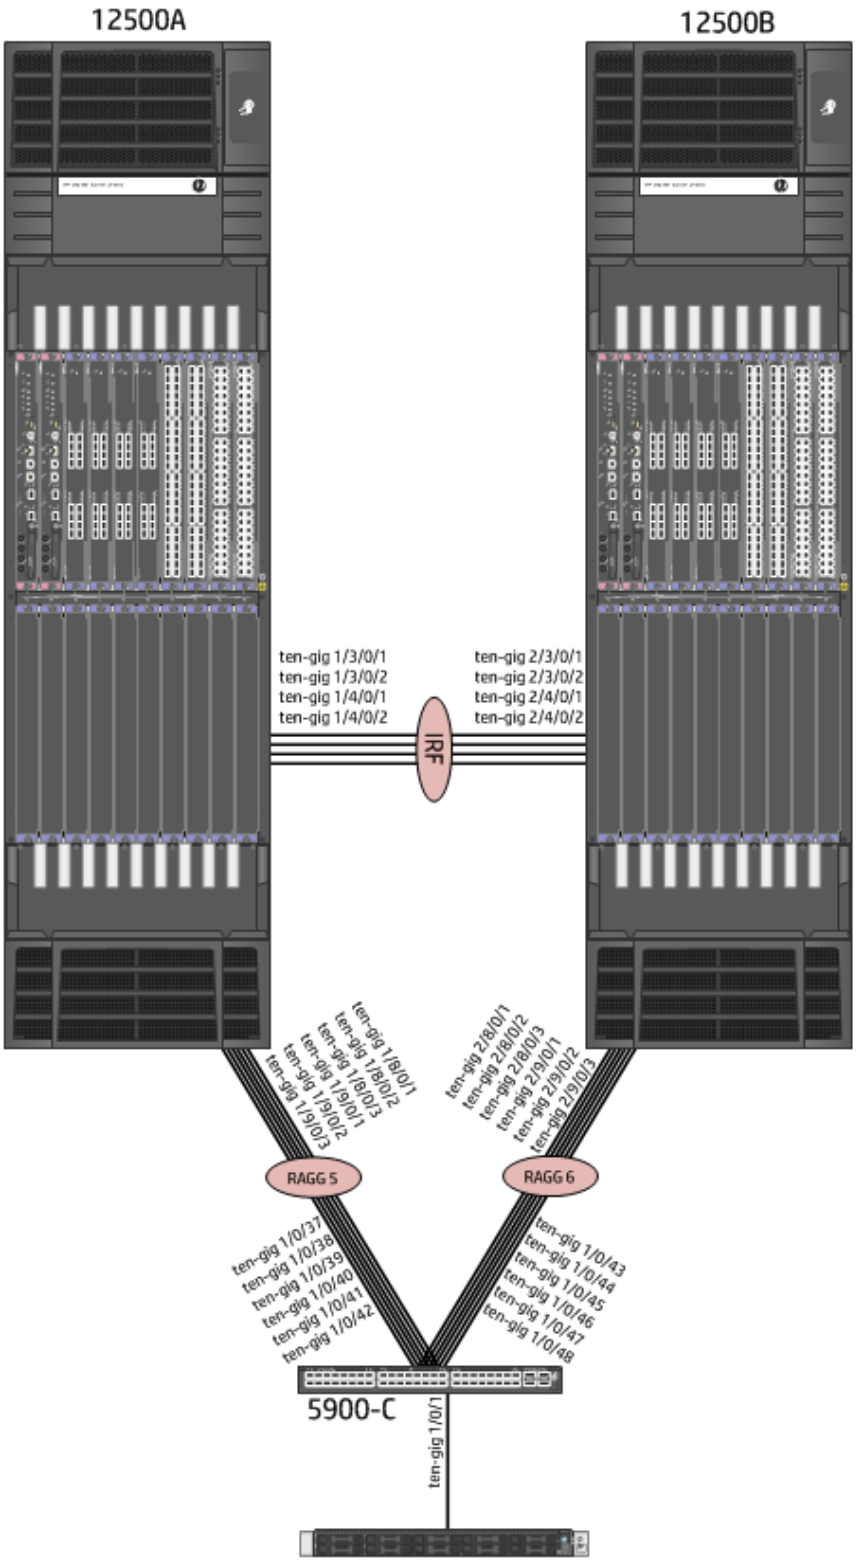

Table 6 2-Tier Layer 3 Design Cabling Information

Link type	Device	Port	Link Type	Remote Device	Port
Route aggregation group 5	12518A	1/8/0/1	10GbE	5900-C	1/0/37
		1/8/0/2			1/0/38
		1/8/0/3			1/0/39
		1/9/0/1			1/0/40
		1/9/0/2			1/0/41
		1/9/0/3			1/0/42
Route aggregation group 6	12518B	2/8/0/1	10GbE	5900-C	1/0/43
		2/8/0/2			1/0/44
		2/8/0/3			1/0/45
		2/9/0/1			1/0/46
		2/9/0/2			1/0/47
		2/9/0/3			1/0/48
Downlink	5900-C	1/0/1	10GbE	Rack Server	LOM1

From a logical view:

Table 7 2-Tier Layer 3 Design VLANs and IP Addresses

Device	VLAN	IP Address	Uplink	Details
12500-A	VLAN301	10.10.301.1 / 24	N/A	Local
	VLAN302	10.10.302.1 / 24		
	VLAN303	10.10.303.1 / 24		
	VLAN304	10.10.304.1 / 24		
	VLAN305	10.10.305.1 / 24		
5900-C	VLAN311	10.10.311.1 / 24	N/A	Local
	VLAN312	10.10.312.1 / 24		
	VLAN313	10.10.313.1 / 24		
	VLAN314	10.10.314.1 / 24		
	VLAN315	10.10.315.1 / 24		

12518A	N/A	5.5.5.1 / 30	Route-aggregation 5	Routed port
12518A	N/A	6.6.6.1 / 30	Route-aggregation 6	Routed port
5900-C	N/A	5.5.5.2 / 30	Route-aggregation 5	Routed port
5900-C	N/A	6.6.6.2 / 30	Route-aggregation 6	Routed port
5900-C	VLAN311	10.10.311.1 / 24	1/0/1	Trunked (tagged)
	VLAN312	10.10.312.1 / 24		
	VLAN313	10.10.313.1 / 24		
	VLAN314	10.10.314.1 / 24		
	VLAN315	10.10.315.1 / 24		

LAN Deployment Procedure

This section describes the configuration steps to deploy the HP 12518 and 5900 platforms for the 2-Tier Layer 3 design.

The LAN procedures listed are in addition to the 12518 procedures listed previously in the 1-Tier BladeServer Design section.

- Rack the 5900 switches and install all necessary physical components to establish a functional HP 5900 (shown in 2-Tier Layer 2 Design section)
- Configure VLANs
  - Create VLAN, names, descriptions
  - Configure VLAN interface IP addresses
  - Undo shutdown VLAN interfaces
  - Configure VLAN MTU, if necessary
- Create and establish uplinks and downlinks
  - Configure Link Aggregation Control Protocol (LACP) for each set of uplinks and downlinks
  - Configure port descriptions to assist in troubleshooting
  - Configure load sharing
- Configure Layer 3 OSPF routing on the 12500s and the 5900s
  - Configure area 0
  - Advertise appropriate networks

LAN Configuration Procedures

Create VLAN-interface and configure its IP address.

Configure 12518A

1) Create VLAN-interface 301 and configure its IP address as 10.10.301.1/24

```
vlan 301
  description To ToR
```

```
  quit
interface vlan-interface 301
  ip address 10.10.301.1 24
  undo shutdown
  quit
```

2) Create VLAN-interface 302 and configure its IP address as 10.10.302.1/24

```
vlan 302
  description To ToR
  quit
interface vlan-interface 302
  ip address 10.10.302.1 24
  undo shutdown
  quit
```

3) Create VLAN-interface 303 and configure its IP address as 10.10.303.1/24

```
vlan 303
  description To ToR
  quit
interface vlan-interface 303
  ip address 10.10.303.1 24
  undo shutdown
  quit
```

4) Create VLAN-interface 304 and configure its IP address as 10.10.304.1/24

```
vlan 304
  description To ToR
  quit
interface vlan-interface 304
  ip address 10.10.304.1 24
  undo shutdown
  quit
```

5) Create VLAN-interface 305 and configure its IP address as 10.10.305.1/24

```
vlan 305
  description To ToR
  quit
interface vlan-interface 305
  ip address 10.10.305.1 24
  undo shutdown
  quit
```

Configure 5900-C

1) Create VLAN-interface 311 and configure its IP address as 10.10.311.1/24

```
vlan 311
  description To Core and server
  quit
interface vlan-interface 311
  ip address 10.10.311.1 24
```

```
  undo shutdown
quit
```

2) Create VLAN-interface 312 and configure its IP address as 10.10.312.1/24

```
vlan 312
  description To Core and server
  quit
interface vlan-interface 312
  ip address 10.10.312.1 24
  undo shutdown
  quit
```

3) Create VLAN-interface 313 and configure its IP address as 10.10.313.1/24

```
vlan 313
  description To Core and server
  quit
interface vlan-interface 313
  ip address 10.10.313.1 24
  undo shutdown
  quit
```

4) Create VLAN-interface 314 and configure its IP address as 10.10.314.1/24

```
vlan 314
  description To Core and server
  quit
interface vlan-interface 314
  ip address 10.10.314.1 24
  undo shutdown
  quit
```

5) Create VLAN-interface 315 and configure its IP address as 10.10.315.1/24

```
vlan 315
  description To Core and server
  quit
interface vlan-interface 315
  ip address 10.10.315.1 24
  undo shutdown
  quit
```

Link Aggregation Configuration Procedure

- Configure Route-Aggregation groups (RAGG 5 and 6) on 12500-A and assign the appropriate ports
- Configure .30 IP address and subnet masks for the RAGG 5 and RAGG 6 aggregate interfaces on 12500-A
- Configure Route-Aggregation groups (RAGG 5 and 6) on 5900-C and assign the appropriate ports
- Configure .30 IP address and subnet masks for RAGG 5 and RAGG 6 aggregate interfaces on 5900-C
- Configure trunk uplink from 5900-C to server carrying tagged VLANs 311 to 315

Configure 12518A

1) Create Layer 3 aggregate interface Route-Aggregation 5 on 12500A, enable MAD, configure an IP address and subnet mask for the aggregate interface, and undo shutdown

```
interface route-aggregation 5
  link-aggregation mode dynamic
  mad enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 1]:
The assigned domain ID is: 1
Info: MAD LACP only enable on dynamic aggregation interface
  quit
  ip address 5.5.5.1 30
  undo shutdown
  quit
```

2) Assign the appropriate Layer 3 interfaces to Route aggregation group 5 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 1/8/0/1
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/8/0/2
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/8/0/3
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/9/0/1
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/9/0/2
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/9/0/3
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
```

3) Create Layer 3 aggregate interface Route-Aggregation 6 on 12500A, enable MAD, configure an IP address and subnet mask for the aggregate interface, and undo shutdown

```
interface route-aggregation 6
  link-aggregation mode dynamic
  mad enable
You need to assign a domain ID (range: 0-4294967295)
[Current domain is: 1]:
The assigned domain ID is: 1
Info: MAD LACP only enable on dynamic aggregation interface
  quit
  ip address 6.6.6.1 30
  undo shutdown
  quit
```

4) Assign the appropriate Layer 3 interfaces to Route aggregation group 6 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 2/8/0/1
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/8/0/2
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/8/0/3
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/9/0/1
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/9/0/2
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 2/9/0/3
  description To 5900-C L3 ToR
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
```

Configure 5900-C

1) Create Layer 3 Route-Aggregation 5 on 5900-C, configure an IP address and subnet mask for the aggregate interface, and undo shutdown

```
interface route-aggregation 5
  link-aggregation mode dynamic
  ip address 5.5.5.2 30
  undo shutdown
  quit
```

2) Assign the appropriate Layer 3 interfaces to link aggregation group 5 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 1/0/37
  description To 12500 l3 Core-A
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/38
  description To 12500 l3 Core-A
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/39
  description To 12500 l3 Core-A
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/40
  description To 12500 l3 Core-A
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/41
  description To 12500 l3 Core-A
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/42
  description To 12500 l3 Core-A
  port link-mode route
  port link-aggregation group 5
  undo shutdown
  quit
```

3) Create Layer 3 Route-Aggregation 6 on 5900-C, configure an IP address and subnet mask for the aggregate interface, and undo shutdown

```
interface route-aggregation 6
  link-aggregation mode dynamic
  ip address 6.6.6.2 30
```

```
  undo shutdown
  quit
```

4) Assign the appropriate Layer 3 interfaces to link aggregation group 6 one at a time, and undo shutdown

```
interface Ten-GigabitEthernet 1/0/43
  description To 12500 l3 Core-B
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/44
  description To 12500 l3 Core-B
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/45
  description To 12500 l3 Core-B
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/46
  description To 12500 l3 Core-B
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/47
  description To 12500 l3 Core-B
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
interface Ten-GigabitEthernet 1/0/48
  description To 12500 l3 Core-B
  port link-mode route
  port link-aggregation group 6
  undo shutdown
  quit
```

5) Configure the 5900 to use either the source and destination IP or the source and destination MAC addresses as the global link-aggregation load sharing criteria. In virtual environments, consider using source/destination IP load-sharing as MAC load-sharing may favor a single link. Both options are shown below.

```
link-aggregation load-sharing mode source-ip destination-ip
link-aggregation load-sharing mode source-mac destination-mac
```

6) Configure Ten-GigabitEthernet 1/0/1 on 5900-C as a trunk port carrying tagged VLANs 311 to 315

```
interface Ten-GigabitEthernet1/0/1
  port link-type trunk
```

```
port trunk permit vlan 311 to 315
undo shutdown
quit
```

Enable OSPF Routing

1) Configure the 12518A with router ID of 1.1.1.1, and configure OSPF Area 0 , and advertise appropriate networks

```
ospf router id 1.1.1.1
  area 0
  network 1.1.1.1 0.0.0.0
  network 5.5.5.0 0.0.0.255
  network 6.6.6.0 0.0.0.255
  network 10.10.0.0 0.0.255.255
  network 172.16.0.0 0.0.255.255
  network 192.168.5.0 0.0.0.255
```

2) Configure the 5900-C with router ID of 2.2.2.2, and configure OSPF Area 0 , and advertise appropriate networks

```
ospf router id 2.2.2.2
  area 0
  stub
  network 2.2.2.2 0.0.0.0
  network 5.5.5.0 0.0.0.255
  network 6.6.6.0 0.0.0.255
  network 10.10.0.0 0.0.255.255
  network 172.16.0.0 0.0.255.255
  network 192.168.5.0 0.0.0.255
```

Rack Server Deployment Procedure Configuration

Since there are many variables when configuring rack servers for specific customer environment, the following items should be considered. Certain setup aspects will vary based on the host OS and are outside the scope of this document.

- iLO management parameters
- Will Intelligent Provisioning be used
- Power options
- NIC options
  - Active/active LACP group
  - Active/Standby NICs

Please refer to the server manuals for your particular server model. An example of the DL380p Gen8 server is linked below.

http://h20000.www2.hp.com/bizsupport/TechSupport/DocumentIndex.jsp?lang=en&cc=us&contentType=SupportManual&prodTypeId=15351&prodSeriesId=5177957

Rack Server Configuration Procedure

Since rack servers are typically custom made to order, the configuration procedure will be different based on the server deployment. Below is a checklist of items that should be performed. A detailed configuration is outside the scope of this document.

- Install components like memory, NICs , hard drives
- Rack the server
- Connect the power supply
- Connect a terminal, or keyboard, monitor and mouse
- Configure iLO settingsDeploy the host OS using either Intelligent provisioning or
- Once host OS is setup configure NIC options
- Configure storage options

Security and network management

HP recommends using only secure management protocols. The examples shown are baseline configurations performed on 12518A running Comware7. Configurations on 5900s will be similar.

Table 8 Security and network management logical

Device	VLAN	IP Address	Uplink	Details
NMS		172.16.1.100		
12518A	VLAN2	172.16.1.1 / 24	1/6/0/3	sFlow agent
12518A		172.16.1.100		sFlow collector

Deployment procedures

- Configure access and security configuration
  - User / SSH / console access
  - Secure VLAN1
- Configure SNMP security, communities, parameters, logging, traps
- Enable LLDP
- Configure sFlow

Configuration procedures

Access and security configuration

Configure User / Telnet / Console

1) Create a local user, enable password control to suppress displaying of local passwords, authorize Telnet, SSH, console, and service-type.

```
password-control enable
local-user admin class manage
  password simple admin
  service-type telnet ssh terminal
  authorization-attribute user-role network-admin
  quit
```

2) Create console user interfaces, enable authentication mode scheme, and assign user role

```
user-interface con 1/0 2/0
  authentication-mode scheme
  user-role network-admin
  quit
```

3) Enable Telnet server, create VTY user interfaces, enable authentication mode scheme, and assign user role

```
telnet server enable
  user-interface vty 0 15
  authentication-mode scheme
  user-role network-admin
  quit
```

Configure SSH

1) Generate the RSA and DSA key pairs on the Telnet sever

```
public-key local create rsa
The range of public key size is (512 ~ 2048).
If the key modulus is greater than 512, it will take a few
minutes.
Press CTRL+C to abort.
Input the modulus length [default = 1024]:
Generating Keys...
.....................++++++
.................++++++
..++++++++
............++++++++

public-key local create dsa
The range of public key size is (512 ~ 2048).
If the key modulus is greater than 512, it will take a few
minutes.
Press CTRL+C to abort.
Input the modulus length [default = 1024]:
Generating Keys...
.+++++++++++++++++++++++++++++++++++++++++++++++++*
........+......+.....+.....................................+
...+................+.........+...+.
```

2) Enable the SSH server function

```
ssh server enable
```

3) Assign an IP address to management interface, which the Telnet client will use as the destination for SSH connection

```
interface management 1/0/0/0
  ip address 192.168.5.1 255.255.255.0
```

Secure VLAN 1

1) Remove all access ports from VLAN1. You can use interface range command to implement the same command to a number of ports at the same time. If the port/s are access port in VLAN 1 only, then they will need to be assigned as an access port to a VLAN other than VLAN 1 – thereby removing VLAN 1 from the ports. The below example assigns ports 1 through 48 as access ports to vlan 2

```
interface range GigabitEthernet 1/0/1 to GigabitEthernet 1/0/48
  port access vlan 2
```

2) If the port/s are trunk ports which permit multiple vlans, including vlan 1 , remove vlan 1 from the trunk configuration

```
interface range GigabitEthernet 1/0/1 to GigabitEthernet 1/0/48
  undo port trunk permit vlan 1
```

3) Delete VLAN1 from bridge-aggregation trunk configurations. The shown configuration does not have VLAN 1 as enabled on the trunks that have been set on the aggregation groups. An example of removing VLAN 1 from a trunk configuration is shown below:

```
interface Bridge-Aggregation 8
  undo port trunk permit vlan 1
```

4) Avoid using VLAN 1 for in-band management. Dedicate a specific VLAN for in-band management to separate user and network control traffic, and assign ports to vlan for management. For the equipment used in this configuration guide, the dedicated management ports should be used.

```
interface vlan-interface 5
  description management-VLAN
  port GigabitEthernet 1/6/0/40
  quit
```

5) Apply an ACL to the management VLAN to restrict only SSH to establish connections. Example below:

```
acl number 3000
  description permit SSH for management only
rule 1 permit tcp source-port ssh destination 129.110.1.2
0.0.0.0
  rule 1 deny tcp source-port any destination 129.110.1.2 0.0.0.0
  quit
traffic classifier ssh_permit
  if-match acl 3000
  quit
traffic behavior any-deny
  filter deny
  quit
qos policy ssh
  classifier ssh_permit behavior any-deny
  quit
interface GigabitEthernet 1/6/0/40
  qos apply policy ssh inbound
  quit
```

Configure SNMP security, communities, parameters, logging, traps.

1) Configure the IP address of the SNMP agent as 172.16.1.100/24 and make sure that the agent and the SNMP Server can reach each other. (Details not shown)
2) Assign the NMS (SNMPv3 group managev3group) read and write access to the objects under the snmp node (OID 1.3.6.1.2.1.11), and deny its access to any other MIB object

```
undo snmp-agent mib-view ViewDefault
snmp-agent mib-view included test snmp
```

```
snmp-agent group v3 managev3group privacy read-view snmp write-
view test
```

3) Add the user managev3user to the SNMPv3 group managev3group, and set the authentication algorithm to MD5, authentication key to authkey, encryption algorithm to DES56, and privacy key to prikey

```
snmp-agent usm-user v3 managev3user managev3group simple
authentication-mode md5 authkey privacy-mode des56 prikey
```

4) Configure contact and physical location information for the agent

```
snmp-agent sys-info contact Mr.IT:1234
snmp-agent sys-info location telephone-closet,3rd-floor
```

5) Enable notifications, specify the NMS at 172.16.1.100 as a trap destination, and set the username to managev3user for the traps

```
snmp-agent trap enable
snmp-agent target-host trap address udp-domain 172.16.1.100
params securityname managev3user v3 privacy
```

6) Configure the SNMP NMS:

- Specify SNMPv3
- Create the SNMPv3 user **managev3user**
- Enable both authentication and privacy functions
- Use MD5 for authentication and DES56 for encryption
- Set the authentication key to authkey and the privacy key to prikey
- Set the timeout time and maximum number of retries
- For information about configuring the NMS, see the NMS manual

Note: The SNMP settings on the agent and the NMS must match.

Enable LLDP

1) Enable LLDP globally (LLDP is enabled on ports by default).

```
lldp enable
```

Configure sFlow

1) Configure the sFlow agent and sFlow collector. Add GigabitEthernet 1/6/0/3 to VLAN 2, and configure the IP address of VLAN-interface 2

```
vlan 2
  port GigabitEthernet 1/6/0/3
  quit
interface vlan-interface 2
  ip address 172.16.1.1 24
  quit
```

2) Specify the IP address of the sFlow agent

```
sflow agent ip 172.16.1.1
```

3) Specify ID 2, IP address 172.16.1.100, the default port number, and description of netserver for the sFlow collector

```
sflow collector 2 ip 172.16.1.100 description netserver
```

4) Set the flow sampling mode and sampling rate. Specify flow collector 2 to receive the flow sampling data

```
interface Gigabitethernet 1/5/0/1
  sflow sampling-mode determine
  sflow sampling-rate 4000
  sflow flow collector 2
  quit
interface Gigabitethernet 2/5/0/1
  sflow sampling-mode determine
  sflow sampling-rate 4000
  sflow flow collector 2
  quit
```

Appendix A

Technologies and Design Recommendations

The following section provides a high level overview of some of the technologies used within the FFRA designs.

LAN

Intelligent Resilient Framework (IRF)

IRF an HP innovation, is a technology that enhances ordinary Ethernet switching designs, allowing substantial improvements in the way Ethernet switches communicate. HP IRF provides the ability to flatten the data center and campus networks, eliminating the need for multiple tiers of aggregation switches and unutilized data paths. IRF provides a framework that enhances Ethernet and provides better link management, utilization, and redundancy.

Why is there a need for IRF?

While STP/RSTP/MSTP are fairly effective in preventing unwanted network loops, convergence can still take several seconds, affecting applications that cannot handle that length of delay. In addition, the performance of STP is poor because it blocks all parallel paths except the one it has selected as active. Even when the network is operating normally, STP can reduce the effective bandwidth (possibly to a degree greater than 50 percent) see figure below.

Figure 44 IRF

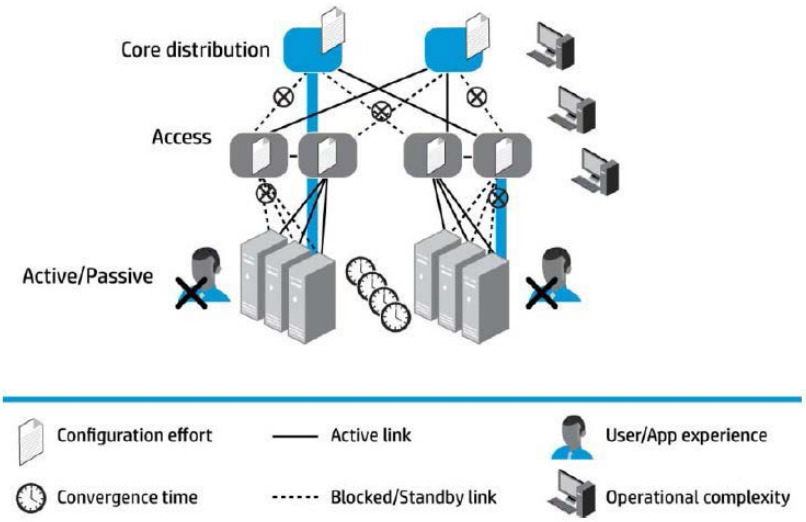

HP IRF can provide a network that is fully resilient, yet also simpler to setup and manage, faster to converge, and easier to scale. IRF simplifies network operations by consolidating management of multiple discrete devices into a single, easy-to- manage virtual switch, in every layer of the network.

Multi-Active Detection (MAD)

If an IRF link failure occurs and members in an IRF system cannot communicate, it could cause an IRF split stack. Two separate IRF systems will be formed. Each system will elect a master and use the IP addresses and

configuration settings assigned to the original IRF system. You can immediately see issues a split stack would cause.

IRF includes a mechanism, called multi active detection (MAD), to quickly discover IRF split stacks. You can implement three types of MADs:

- LACP MAD
- BFD MAD
- ARP MAD

IRF design recommendations:

- Utilize at least 2 physical ports for each IRF logical port for resiliency; however, a best practice is to have at least two IRF links between each 12500 and when using 4 chassis IRF to utilize a ring topology
- Utilize IRF ports from separate modules within each chassis for module resiliency
- Configure MAD detection method – best practice is to use both BFD MAD and LACP MAD, but if no Comware based switch is available for LACP MAD then use BFD MAD
- Before establishing an IRF fabric, make sure that the system working mode of the member switches is the same. If not, the IRF fabric cannot be established
- The member switches of an IRF fabric must work in the same rule match mode. This means that you must configure the acl ipv6 enable command, or the acl ipv6 disable command on the switches
- The member switches of an IRF fabric must be configured with the same VPN label processing mode.
- Before establishing an IRF fabric, check that enhanced IRF mode is enabled on all member switches or disabled on all member switches. If enhanced IRF mode is enabled on some member switches but disabled on the others, the IRF fabric cannot be established

VLANs

Ethernet is a network technology based on the CSMA/CD mechanism. As the medium is shared, collisions and excessive broadcasts are common on Ethernet networks. To address the issue, VLAN was introduced to break a LAN down into separate VLANs. VLANs are isolated from each other at Layer 2. A VLAN is a bridging domain, and all broadcast traffic is contained within it. A VLAN is logically divided on an organizational basis rather than on a physical basis. For example, all workstations and servers used by a particular workgroup can be assigned to the same VLAN, regardless of their physical locations.

VLAN and IP Design Recommendations

VLAN and IP design recommendations and considerations:

- Utilize port based VLANs
- Use trunk links to blade switches so they can carry multiple VLANs
- Make sure to secure VLAN1 (see access and security configuration section)

Ethernet link aggregation

Ethernet link aggregation, or simply link aggregation, combines multiple physical Ethernet ports into one logical link, called an aggregate link. Link aggregation delivers the following benefits:

- Increases bandwidth beyond the limits of any single link. In an aggregate link, traffic is distributed across the member ports

- Improves link reliability. The member ports dynamically back up one another. When a member port fails, its traffic is automatically switched to other member ports

Ethernet link aggregation Design Recommendations

- Use dynamic Link Aggregation Control Protocol (LACP)
- Enable broadcast storm suppression characteristics in order to reduce operating risks
- Be sure to be aware of any link aggregation limits the device used may have. At the time of release of this document the 12500 supports:
    - Up to 240 aggregation groups
    - A single 12500 chassis supports up to 12 selected ports
    - Multiple 12500 chassis using IRF with CW5 (R1825) or CW7, now support 12xN in a link aggregation group, with N being the number of chassis in the IRF domain

IGP routing (OSPF)

OSPF, as an Interior Gateway Protocol (IGP), functions between routing devices within the same domain, or autonomous system (AS), which is generally defined as a group of devices working under the control of the same entity. It enables the routing devices to exchange information about each other about the IP subnets within the AS to discover routes between them.

IGP routing (OSPF) Design Recommendations

- This scenario described below in the L3 ToR section will utilize a single OSPF area, however in larger more complex OSPF deployments more areas and route summaries may be recommended to reduce routing table size. These types of deployments are out of scope of this document
- Additionally, the example below configures the L3 ToR switches as a stub area. The core ABR injects a default route into the stub area. This ensures that routers in the stub area will be able to route traffic to external destinations without having to maintain all of the individual external routes

SNMP

SNMP is an Internet standard protocol widely used by management station so they can access and operate the devices on a network, regardless of their vendors, physical characteristics, and interconnect technologies.

SNMP enables network administrators to read and set the variables on managed devices for state monitoring, troubleshooting, statistics collection, and other management purposes.

SNMP Design Recommendations

- HP recommends using SNMPv3 which uses a user-based security model (USM) to secure SNMP communication. You should configure authentication and privacy mechanisms to authenticate and encrypt SNMP packets for integrity, authenticity, and confidentiality

Link Layer Discovery Protocol (LLDP)

The LLDP protocol operates on the data link layer to exchange device information between directly connected devices. With LLDP, a device sends local device information (including its major functions, management IP address, device ID, and port ID) as TLV (type, length, and value) triplets in LLDP Data Units (LLDPDUs) to the

directly connected devices, and at the same time, stores the device information received in LLDPDUs sent from the LLDP neighbors in a standard management information base (MIB). It allows a network management system to quickly detect and identify Layer 2 network topology changes.

LLDP Design Recommendations

- LLDP should be enabled on all interfaces
- To make your switch work with Cisco devices, you should enable CDP compatibility

Access and security configuration

Telnet / SSH

User interface (also called line) allows you to manage and monitor sessions between the terminal and switch when you are using the console port, AUX port, and asynchronous serial interfaces to log in to the switch by Telnet or SSH.

One user interface corresponds to one user interface view where you can configure a set of parameters, such as whether to authenticate users at login, whether to redirect the requests to another device, and the user level after login. When the user logs in through a user interface, the connection follows these parameter settings, thus implementing centralized management of various sessions.

Secure Shell (SSH) is a network security protocol. Using encryption and authentication, SSH can implement secure remote access and file transfer over an insecure network. The 12500 switch supports both secure Telnet and secure FTP.

Telnet/SSH design recommendations:

- HP recommends using SSH only to access the switch.
- Generate DSA and RSA key pairs
- Enable SSH and SFTP
- Configure the user interfaces
- Configure the clients host public key, and SSH user

Secure VLAN 1 Design Recommendations

- Remove all access ports from VLAN1
- Delete VLAN1 from trunk configurations
- Avoid using VLAN 1 for in-band management. Dedicate a specific VLAN for in-band management to separate user and network control traffic
- Apply an ACL to the management VLAN to restrict only interested protocols such as SSH or SSH File Transfer Protocol (SFTP) to establish connections. Additionally apply a Quality of Service (QoS) ACL to rate limit the number of ping traffic allowed

sFlow

Sampled Flow (sFlow) is a traffic monitoring technology mainly used to collect and analyze traffic statistics. The sFlow system involves an sFlow agent embedded in a switch and a remote sFlow collector. The sFlow agent collects traffic statistics and packet information from the sFlow-enabled interfaces on the switch, encapsulates them into sFlow packets. When an sFlow packet buffer overflows, or an sFlow packet ages out (the aging time is one second), the sFlow agent sends the packet to a specified sFlow collector. The sFlow collector analyzes the sFlow packets and displays the results.

sFlow has the following two sampling mechanisms:

- Flow sampling: Packet-based sampling, used to obtain packet content information
- Counter sampling: Time-based sampling, used to obtain port traffic statistics

Virtual Connect (VC)

Virtual Connect server profiles

The server profile, an essential component to configuring VC modules, provides a link between the server and the networks and fabrics defined in VC. The server profile can include MAC and WWN addresses, as well as boot parameters for the various connection protocols supported by VC. After being defined, the server profile can be assigned to any server blade within the Virtual Connect domain. VCM supports up to 256 profiles within the domain.

A Virtual Connect server profile consists of connections that group attributes related to server connectivity for the various protocols supported by Virtual Connect modules. These protocols are Ethernet, iSCSI, Fibre Channel over Ethernet (FCoE), and Fibre Channel.

Server profiles are associated with a specific enclosure device bay. After a profile is assigned, the Virtual Connect Manager configures the server blade in that device bay with the appropriate MAC, PXE, WWN, and SAN boot settings and connects the appropriate networks and fabrics.

Shared Uplink Sets (SUS)

The SUS provides the ability to support VLAN tagging and frame forwarding based on the VLAN tags of those frames. The SUS connects one or many server NICs to one or many uplink ports. A SUS should be configured for the specific VLANs it will support. If support for additional VLANs is required, those VLANs need to be configured within the SUS.

When connecting a server NIC to a network within a SUS, there are two choices provided. The key difference between these two options is the state in which the frame is passed to the server NIC. When configuring a server NIC for network connections within the server profile, choose from the following methods;

1. Selecting a single network – which would be mapped to a specific VLAN. If a single network is selected, the frames will be presented to the server NIC WITHOUT a VLAN tag. In this case the host operating system does not need to understand which VLAN it resides in. When the server transmits frames back to VC, those frames will not be tagged, however; Virtual Connect will add the VLAN tag and forward the frame onto the correct VLAN
2. Selecting multiple networks (Map VLAN Tags)– which would provide connectivity to several VLANs. The Map VLAN Tags feature provides the ability to use a Shared Uplink Set to present multiple networks to a single NIC. If you select Multiple Networks when assigning a Network to a server NIC, you will have the ability to configure multiple Networks (VLANS) on that server NIC. At this point VC tags ALL the packets presented to the NIC — unless the Native check box is selected for one of the networks, in which case packets from this network (VLAN) will be untagged, and any untagged packets leaving the server will be placed on this Network (VLAN)

With Mapped VLAN Tags, you can create a Shared Uplink Set that contains ALL the VLANs you want to present to your servers, then present only ONE network (the one associated with the VLAN we want the server NIC in) to the

Windows, LINUX or the ESX Console NIC, then select Multiple Networks for the NIC connected to the ESX vSwitch and select ALL the networks that we want presented to the ESX host vSwitch. The vSwitch will then break out the VLANs into port groups and present them to the guests. Using Mapped VLAN Tags minimizes the number of uplinks required.

A Shared Uplink Set can be configured to support both tagged and un-tagged network traffic to a server NIC, which simplifies the overall configuration and minimizes the number of uplink cables required to support the network connections.

FlexNIC

HP FlexFabric technology is exclusive to HP Virtual Connect environments. When a server's 10GbE capable FlexFabric adapter is connected to an HP Virtual Connect Flex-10/10D or FlexFabric Module, each NIC port becomes four individual NICs, called FlexNICs. Although these four FlexNICs share a single 10Gb physical interface, Virtual Connect is able to keep traffic for the FlexNICs isolated, and each FlexNIC is assigned to one or more distinct Virtual Connect networks. When connected to a VC FlexFabric module, the 2nd FlexNIC from each physical NIC can be dedicated to either FCoE or iSCSI traffic.

Each FlexNIC can be assigned a different transmit bandwidth (from 100Mb to 10Gb), which is enforced by hardware mechanisms. The FlexNICs share a total of 10Gb, so one could be set to 5Gb, one could be set to 1Gb, and the remaining two could each be set to 2Gb. The total shareable bandwidth cannot exceed 10Gb.

Figure 45 FlexNIC and NIC Mapping

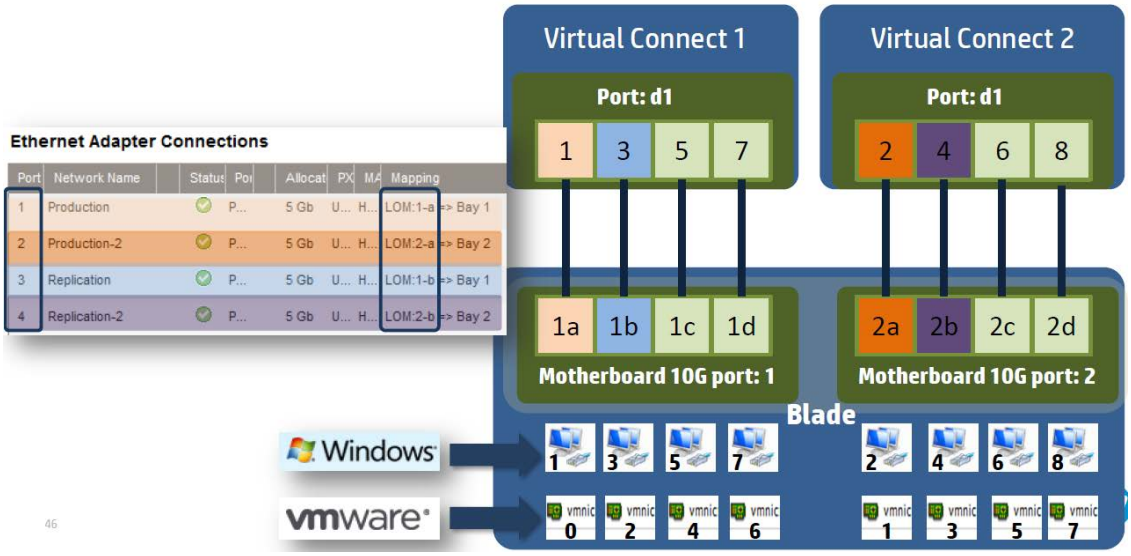

FlexNIC important notes:

- NC553i FlexFabric adapters support bandwidth range from 100Mb ⇨ 10Gb
- NC551i/m FlexFabric adapters support bandwidth range from 1Gb ⇨ 10Gb

Figure 46 FlexNIC bandwidth assignment example

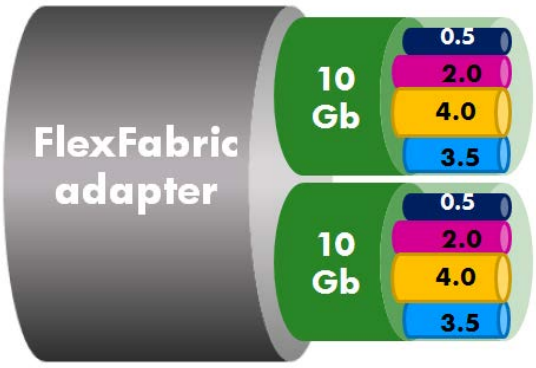

Figure 47 The power of the FlexNIC

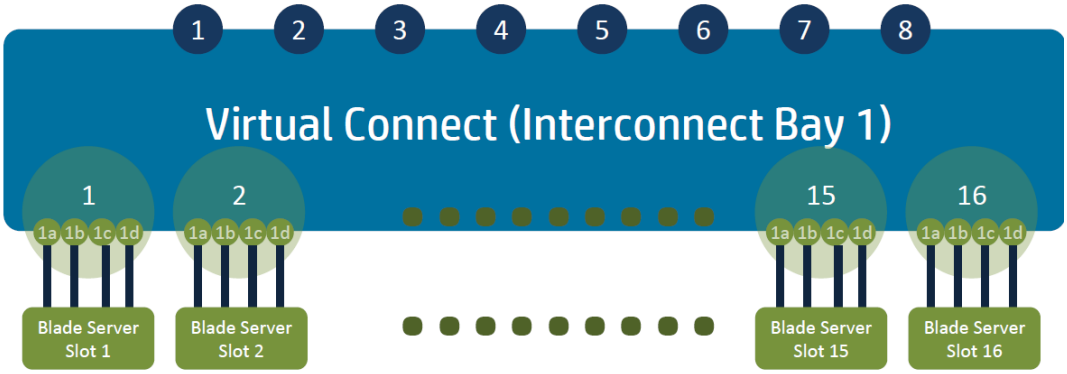

One Flex-10 or FlexFabric offers up to 64 downlinks ports to 16 blade servers

A pair of Flex-10 or FlexFabric offers up to 128 downlinks ports to 16 blade servers

Virtual Connect SAN Fabrics

Beginning with Virtual Connect 3.70, there are two supported VC SAN fabric types; FabricAttach fabrics and DirectAttach fabrics.

FabricAttach VC SAN fabrics

The VC-Fibre Channel (FC) and FlexFabric modules enable the c-Class administrator to reduce FC cabling by using N_Port_ID virtualization (NPIV). The HP VC-FC and FlexFabric modules act as an FC connectivity aggregator, where each NPIV-enabled N-port uplink can carry the FC traffic for multiple HBAs or FlexFabric adapters.

Because the uplink ports for VC-FC and FlexFabric modules are N-ports, the modules can be connected to any data center Brocade, McData, Cisco, and Qlogic FC switch that supports the NPIV protocol. When the server blade HBAs or FlexFabric adapters log in to the fabric through the VC-FC or FlexFabric modules, the adapter WWN is visible to the FC switch name server and can be managed as if it was connected directly.

Figure 48 FlexFabric and FC-VC FabricAttach SAN fabrics

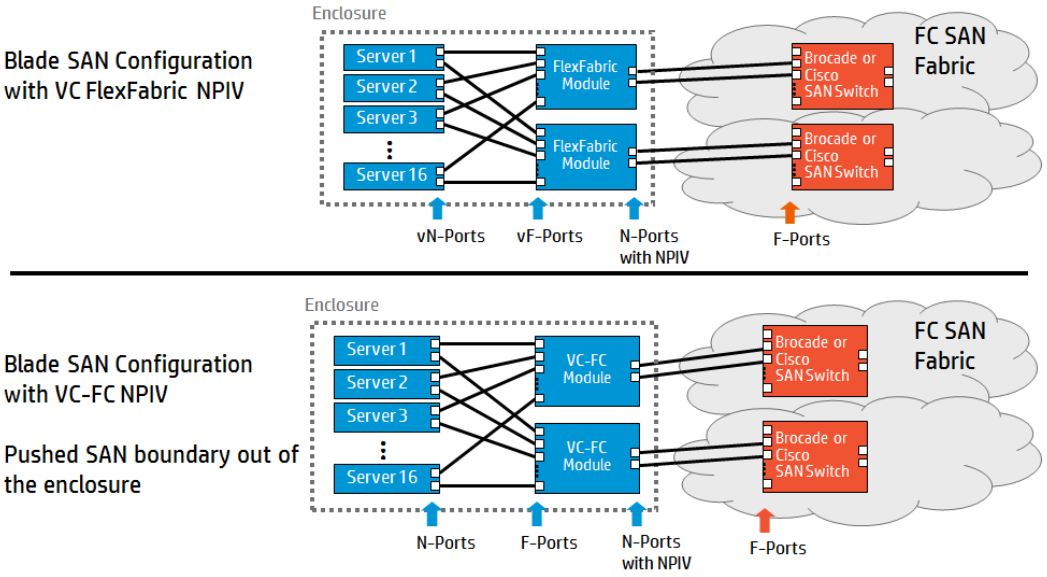

FabricAttach VC SAN important notes

- By default, all of the VC-FC module uplinks are grouped into a single fabric, distributing connectivity from all server blades in the enclosure
- By default, all of the FlexFabric FC-capable uplinks are configured as Ethernet until they are configured as part of the VC SAN fabric. After the FC-capable uplinks are configured as part of the VC SAN fabric, the FC SFP transceivers connected to those uplinks become enabled and allow connectivity to the data center SAN fabric
- To create a proper Virtual Connect fabric, all VC-FC or FlexFabric module uplinks that are included in the fabric must be connected to the same SAN fabric as shown in the following figure

Figure 49 VC-FC or FlexFabric module uplink example

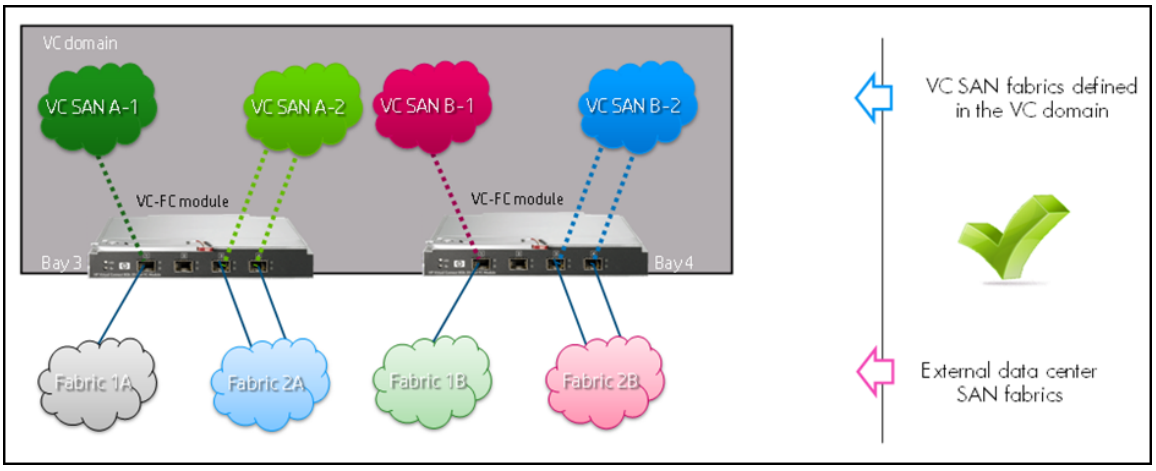

DirectAttach VC SAN fabrics

Virtual Connect Direct-Attach Fibre Channel for 3PAR Storage Systems transforms the efficiency of server and storage connectivity by eliminating the need for complex, multi-tier SANs.

DirectAttach fabric requires HP VC FlexFabric modules and is supported on a subset of 3PAR Storage Systems (check the 3PAR documentation for up to date supported list). When the uplink ports of a FlexFabric module are configured for a DirectAttach fabric, the uplink ports employ simplified SAN fabric services combined with auto-configured initiator-based zoning. This allows the supported storage systems to be directly attached to the uplink ports on a module without the need for an intermediate SAN fabric. The servers and the supported storage devices log in to the VC SAN fabric independently.

Figure 50 FlexFabric DirectAttach FC SAN fabrics

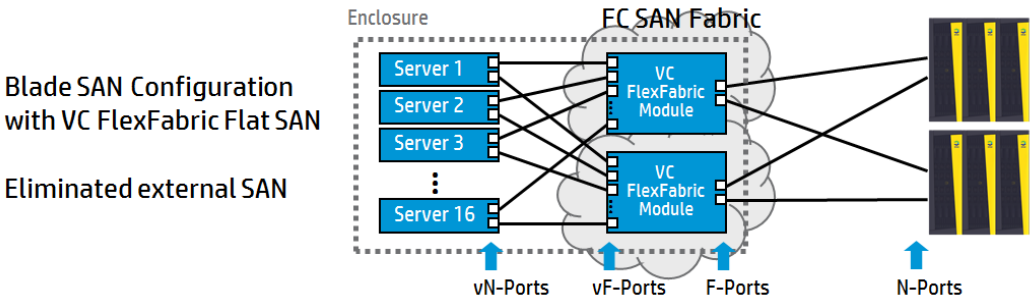

DirectAttach VC SAN important notes

- The DirectAttach fabric is only supported with the HP VC FlexFabric 10Gb/24-port Module when it is connected to one or more supported HP 3PAR storage systems
  - The minimum required version of HP Virtual Connect firmware is v3.70
  - The minimum required version of HP 3PAR InForm OS is v3.1.1 MU1
- The following storage systems are not supported: HP MSA/EVA/XP storage systems, HP StoreOnce Backup appliance, HP LeftHand Storage systems, HP Tape and Virtual Tape Libraries, and any third-party storage solution.
- When creating the DirectAttach fabric, all participating uplinks can be connected to the same 3PAR storage system in order to form a VC SAN fabric correctly.
- For more control over the uplink port utilization, you can create several DirectAttach VC SAN fabrics connected to the same 3PAR storage system. This configuration can assist the distribution of servers according to server I/O needs and workloads.

Mixed FabricAttach and DirectAttach VC SAN fabrics

Mixing FabricAttach and DirectAttach VC SAN fabrics is fully supported in the same Virtual Connect domain. This scenario can be useful if you need to attach additional storage systems that are not supported today with the DirectAttach fabrics.

Virtual Connect and IRF Cabling

There is a common misunderstanding people tend to have when connecting Virtual Connect with IRF configured switches. The first example presented below is the recommended solution.

In this configuration, one Shared Uplink Set will be configured on each Virtual Connect module (two total). The logical IRF switch also has two link aggregation groups (LACP) configured to peer with the Virtual Connect SUS. This configuration provides an Active/Active Virtual Connect design. Note that more ports can be added to increase uplink bandwidth, if necessary.

For more information on Active/Standby and Active/Active designs, see the HP Virtual Connect FlexFabric Cookbook.

Figure 51 Recommended IRF to VC FlexFabric Cabling

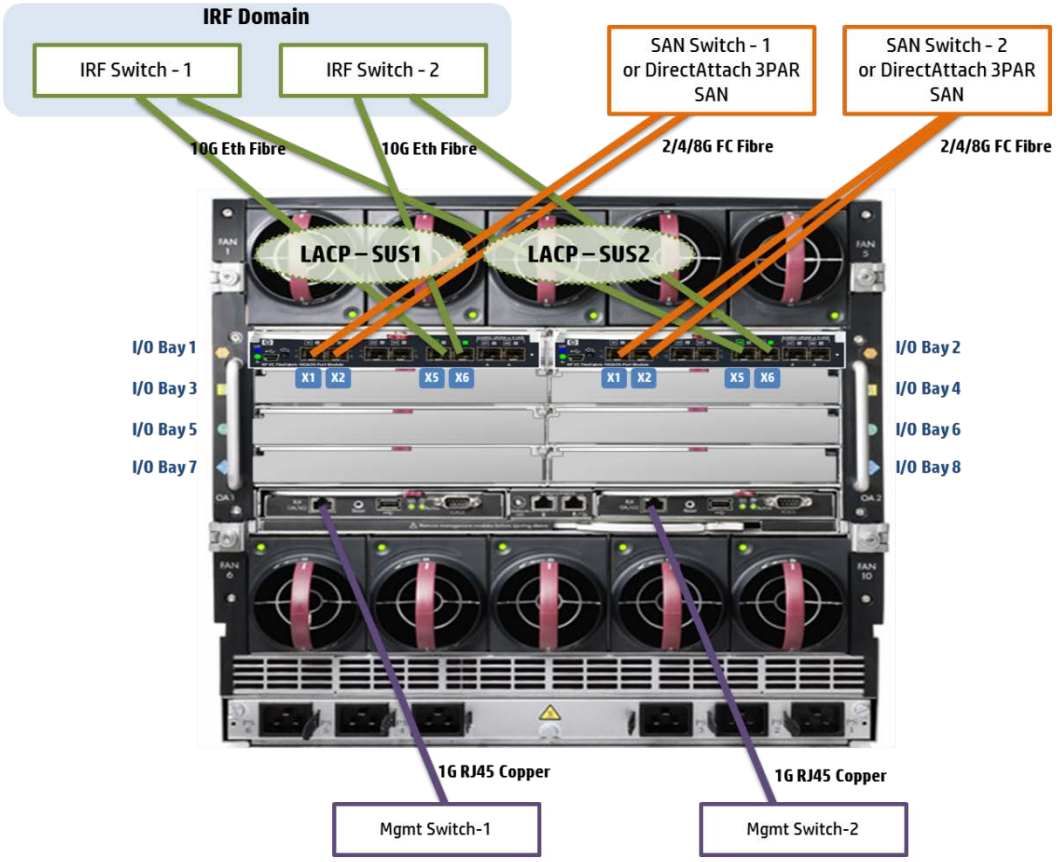

3PAR

HP 3PAR StoreServ Storage concepts and terminology

HP 3PAR StoreServ includes both hardware components that physically store data and software applications that manage data. For more information about hardware platforms, see the section titled “HP 3PAR StoreServ Hardware” in this document. For more information about system software applications and features, see the section titled “HP 3PAR OS Software.”

HP 3PAR StoreServ is comprised of the following logical data layers:

- Physical disks
- Chunklets
- Logical disks
- Common provisioning groups
- Virtual volumes

The relationship between system data layers is illustrated below. Each layer is created from elements of the layer above. Chunklets are drawn from physical disks. Logical disks are created from groups of chunklets. Common provisioning groups (CPGs) are groups of logical disks. And virtual volumes use storage space provided by CPGs. The virtual volumes are exported to hosts and are the only data layer visible to hosts.

Figure 52 3PAR logical data layers

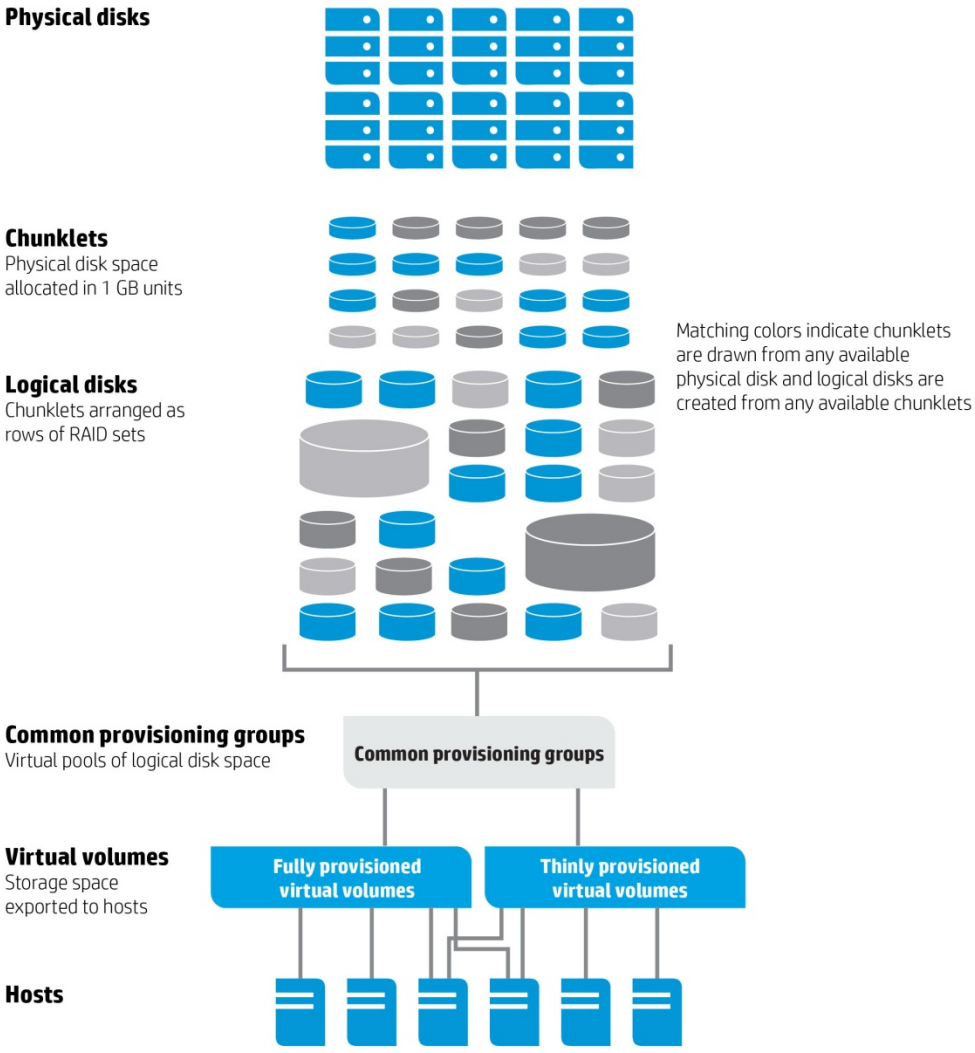

Physical disks

A *physical disk* is a hard drive mounted on a drive magazine located in an HP 3PAR StoreServ drive enclosure.

Chunklets

Physical disks are divided into *chunklets*. Each chunklet occupies contiguous space on a physical disk. On F-Class and T-Class systems, all chunklets are 256 MB. On HP 3PAR StoreServ 10000 and 7000, all chunklets are 1 GB. Chunklets are automatically created by the HP 3PAR OS and they are used to create logical disks. A chunklet is assigned to only one logical disk.

Logical disks

A *logical disk* is a collection of chunklets arranged as rows of RAID sets. Each RAID set is made up of chunklets from different physical disks. Logical disks are pooled together in common provisioning groups, which allocate space to virtual volumes. The underlying logical disks are automatically created by the HP 3PAR OS when you create VVs. The RAID type, space allocation, growth increments, and other logical disk parameters are specified when you create a CPG or can be modified later. HP 3PAR StoreServ support the following RAID types:

- RAID 10 (RAID 1)
- RAID 50 (RAID 5)
- RAID Multi-Parity (MP) or RAID 6

Common provisioning groups (CPGs)

A *CPG* is a virtual pool of logical disks that allocates space to virtual volumes on demand. A CPG allows up to 4,095 virtual volumes to share the CPG's resources. You can create fully provisioned virtual volumes and thinly provisioned virtual volumes (TPVVs) that draw space from a CPG's logical disk pool. It is important to note that if no volumes (thick or thin) have been created in a CPG, it consumes no space.

Virtual volumes (VVs)

Virtual volumes draw their resources from CPGs and are exported as Logical Unit Numbers (LUNs) to hosts. Virtual volumes are the only data layer visible to the hosts. You can create physical copies or virtual copy snapshots of virtual volumes. Full copies remain available if the original base volume becomes unavailable. Before creating virtual volumes, you must first create CPGs to allocate space to the virtual volumes.

Fully provisioned virtual volumes

A *fully provisioned* virtual volume is a volume that uses logical disks that belong to a CPG. Unlike TPVVs, fully provisioned virtual volumes have a set amount of user space that is allocated for user data. The fully provisioned volume size is fixed, and the size limits range from 1 GB to 16 TB.

Thinly provisioned virtual volumes

A *thinly provisioned* virtual volume (TPVV) is a volume that uses logical disks that belong to a CPG. TPVVs associated with the same CPG draw space from that pool as needed, allocating space on demand in small increments for each TPVV. As the volumes that draw space from the CPG require additional storage, the HP 3PAR OS automatically creates additional logical disks and adds them to the pool until the CPG reaches the user-defined growth limit which restricts the CPG's maximum size.

Physical copies

A *physical copy* duplicates all the data from a base volume to a destination volume. The base volume is the original volume that is copied to the destination volume. The physical copy on the destination volume remains available if the original base volume becomes unavailable. Unlike a virtual copy or snapshot, a physical copy can maintain the performance of the base virtual volume provided the physical copy has the same disk characteristics (type, speed, RAID level, etc.).

In addition, the destination volume must have a user space size at least as large as the user space of the base volume being copied. In addition, the HP 3PAR OS now allows the export of the physical copy immediately after the creation of the copy, while the data copy continues to completion in the background.

Note: With an HP 3PAR Remote Copy Software license, physical copies can be copied from one HP 3PAR StoreServ system to another using Remote Copy. For additional information, see the *HP 3PAR Remote Copy User's Guide*.

Virtual copy snapshots

A *snapshot* is a virtual copy of a base volume. The base volume is the original volume that is copied. Unlike a physical copy, which is a duplicate of an entire volume, a virtual copy only records changes to the base volume. This allows an earlier state of the original virtual volume to be recreated by starting with its current state and rolling back all the changes that have been made since the virtual copy was created.

You can make snapshots of fully provisioned virtual volumes, TPVVs, physical copies, or another virtual copy snapshot. Snapshots are created using copy-on-write techniques available only with the HP 3PAR Virtual Copy Software license. Thousands of snapshots of each virtual volume can be created assuming that there is sufficient storage space available.
It is worth noting that snapshots do not consume any space unless data on the base volume has been updated and the original data copied to the Snapshot Data Space. Changed data is copied only once regardless of the number of snapshots taken.

Note: Creating virtual copies requires an HP 3PAR Virtual Copy Software license.

Exporting virtual volumes

For a host to see a virtual volume, the volume must be exported as a logical unit number (LUN). Volumes are exported by creating virtual volume-LUN pairings (VLUNs) on the system. When you create VLUNs, the system produces both *VLUN templates* that establish export rules, and *active VLUNs* that the host sees as LUNs or attached disk devices.

HP 3PAR Operating System Software Suite

HP 3PAR 7000 Operating System Software Suite is the foundation software of HP 3PAR StoreServ 7000 Storage, combining advanced virtualization capabilities with simple storage management, high efficiency, and world class performance. The included comprehensive thin provisioning capabilities allow your storage to start thin, get thin and stay thin. System Tuner and Autonomic Rebalance help maintain high performance over time. Migrating your existing data from HP EVA and 3PAR systems to HP 3PAR StoreServ 7000 is easy with the included 180 day Online Import license. HP 3PAR SmartStart Software, included in the suite, guides you through the configuration of the service processor, StoreServ Storage and the application hosts - making storage setup virtually effortless. HP 3PAR 7000 Operating System Software Suite is a required purchase.

The Operating System Suite includes the following functionality and features:

- HP 3PAR Thin Provisioning
- HP 3PAR Thin Conversion
- HP 3PAR Thin Persistence
- HP 3PAR Thin Copy Reclamation
- HP 3PAR Autonomic Rebalance
- HP 3PAR System Tuner
- HP 3PAR Management Console

- EVA to 3PAR Online Import**
- HP 3PAR Host Explorer
- HP 3PAR SmartStart
- HP 3PAR Virtual Service Processor
- HP 3PAR Multipath Software for Windows 2003

For more information see http://h18006.www1.hp.com/storage/software/3par7000/oss/index.html

For more information

To read more about HP Networking Products, go to http://www.hp.com/go/networking

Get connected

hp.com/go/getconnected

Current HP driver, support, and security alerts

delivered directly to your desktop

© Copyright 2013 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.3

Created April 2013; Rev. 1.0