

Migrating Storage on HP-UX Servers To MDS Switches

Version 1.0

Venkat Kirishnamurthy

January 22, 2004

Table of Contents

1	OVERVIEW	7
2	CONCEPTS	9
2.1	FIBRE CHANNEL IDENTIFIER [FC_ID].....	9
2.2	VOLUME SET ADDRESSING (VSA).	9
2.3	LOGICAL VOLUME MANAGER (LVM)	11
2.3.1	LVM basics.....	11
2.3.2	Physical Volume, Volume Group and Logical Volumes	13
2.3.3	LVM configuration file (/etc/lvmtab).....	16
2.3.4	LVM limitations	16
2.4	MULTI PATH ON HP-UX	16
2.5	DISK ADDRESSING, HP-UX AND FIBRE CHANNEL HBAS.	17
3	NON DISRUPTIVE MIGRATION – 1	21
3.1	MIGRATING STORAGE PORT INDEPENDENT OF ITS FC_ID	21
3.1.1	Migration Steps.....	22
3.1.2	Step 1: Deactivate one of the redundant paths.....	22
3.1.3	Step 2: Move the disconnected path MDS.....	28
3.1.4	Step 3: Scan and list the current status of the devices	28
3.1.5	Step 4: Replace the other Brocade to a full MDS fabric.	34
3.2	PROS AND CONS	35
3.2.1	Pros.....	35
3.2.2	Cons	35
4	NON DISRUPTIVE MIGRATION – 2 (USING THE PERSISTENT FC_ID FEATURE).....	36
4.1	FC_ID REMAINS PERSISTENT ACROSS THE SWITCHES (BEFORE AND AFTER MIGRATION).	36
4.2	MIGRATING THE PORTS ON BROCADE TO MDS.....	37
4.2.1	Gather data before removing the path through Brocade.	37
4.2.2	Connect path through MDS1	40
4.2.3	MDS1 configuration steps.....	41
4.3	MIGRATING THE PORTS ON McDATA TO MDS	41
4.3.1	Gather data before removing the path through McData switch	42
4.3.2	Remove the path through McData switch	45
4.3.3	Connect path through MDS2	45
4.4	PROS AND CONS.....	48
4.4.1	Pros.....	48
4.4.2	Cons	48
5	DISRUPTIVE MIGRATION	49
5.1	MIGRATION: USING LVM IMPORT AND EXPORT OF VOLUME GROUPS.	49
5.1.1	Gather volume group information.	49
5.1.2	Create export map files for the affected volume groups.....	53
5.1.3	Rescan the devices and import the volume group	55
5.2	PROS AND CONS	58
5.2.1	Pros.....	58
5.2.2	Cons	58
6	APPENDIX A	59
6.1	FC_ID, BROCADE AND McDATA SWITCHES AND HP-UX HARDWARE PATH	59
6.1.1	Examples: Hardware path through Brocade switches.	59
6.1.2	Examples: Hardware path through McData switches.	60
7	APPENDIX B.....	62

7.1	VERITAS VOLUME MANAGER	62
7.1.1	VxVM basics.....	62
7.1.2	Multi path with Veritas Volume Manager.....	64

Table of Figures

FIGURE 1: TOPOLOGY OVERVIEW	7
FIGURE 2: 24 BIT FIBRE CHANNEL IDENTIFIER	9
FIGURE 3: PARALLEL SCSI DEVICE ADDRESSING.....	9
FIGURE 4 : VOLUME SET ADDRESSING ON ARRAYS.....	10
FIGURE 5: 14 BIT VSA ADDRESS	11
FIGURE 6: LVM LAYOUT ON A PHYSICAL VOLUME	12
FIGURE 7: VOLUME GROUP WITH OUT PV LINKS	13
FIGURE 8: VOLUME GROUP WITH PV LINKS	13
FIGURE 9: LVM VOLUME CREATION	14
FIGURE 10: DEVICE ADDRESSING ON HP-UX SERVERS.....	17
FIGURE 11: TOPOLOGY WITH HARDWARE PATHS.....	18
FIGURE 12: IOSCAN -FNKC DISK OUTPUT	19
FIGURE 13: HARDWARE PATH'S RELATIONSHIP TO FC_ID ON HP-UX	20
FIGURE 14: FCID AND THE HARDWARE PATH RELATION SHIP	20
FIGURE 15: INITIAL CONFIGURATION	21
FIGURE 16: REMOVE THE PATH THROUGH BROCADE.....	26
FIGURE 17: MDS SWITCH IN THE FABRIC	28
FIGURE 18: ALL MDS TOPOLOGY	34
FIGURE 19: TOPOLOGY WITH BROCADE AND McDATA SWITCHES	36
FIGURE 20: PATH THROUGH BROCADE IS DISCONNECTED.	39
FIGURE 21: MDS1 IN THE FABRIC	40
FIGURE 22 : DISCONNECT McDATA SWITCH FORM TOPOLOGY.....	45
FIGURE 23: ALL MDS TOPOLOGY (MIGRATION 2)	46
FIGURE 24: TOPOLOGY FOR DISRUPTIVE MIGRATION	49
FIGURE 25 : DISCONNECT THE CABLES FROM THE BROCADE SWITCHES	55
FIGURE 26: ALL MDS TOPOLOGY AFTER MIGRATION.....	55
FIGURE 27 : FC_ID FROM THE BROCADE EXAMPLE	59
FIGURE 28 : FC_ID FROM THE McDATA EXAMPLE	61
FIGURE 29: VxVM DISK LAYOUT	62
FIGURE 30: VxVM VOLUME CREATION	63

Revision History

Date	Author	Version	Comments
09/10/2003	Venkat Kirishnamurthy	0.1	Initial text
10/14/2003	Venkat Kirishnamurthy	0.2	modification based on inputs from Asim
10/30/2003	Venkat Kirishnamurthy	0.5	modifications based on inputs from Pierre
1/21/2004	Venkat Kirishnamurthy	1.0	Rev 1.0 of the doc.

Preface

This document provides a reference guide to help in the migration of HP-UX servers from a SAN built with non MDS switches to a SAN built using MDS switches.

Audience

This reference document is designed for use by Cisco TAC, Sales, Support Engineers, Professional Service Partners, Systems Administrators and others, responsible for the design and deployment of Storage Area Networks in the data center environment.

Inquiries

Please send your comments via email to:

Email: kvenkat@cisco.com

State the document name, page number, and details of the request. Thank you.

CCIP, the Cisco Arrow logo, the Cisco *Powered* Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ Logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, LightStream, MGX, MICA, the Networkers logo, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratum, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0208R)

Migrating Storage on HP-UX Servers To MDS Switches

Copyright (c) 2003, Cisco Systems, Inc.

All rights reserved.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS. THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY. The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California. NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Introduction

The focus of this document is to discuss details about the various options/methods available for migrating SANs with HP-UX servers from 3rd party SAN switches (Brocade, McData etc.) to MDS switches. The process of migration and its impact on downtime varies with methods used for migration and applications running on the server.

The storage/system administrators can then decide which method best suits the business needs and use that method for migration.

1 Overview

The topology used in this paper is shown in Figure 1. The topology consists of a storage array connected to a HP-UX system through two SAN switches. The host SJHP2 is connected to switch 1 through Host Bus Adapter (HBA) in slot 0/1/1/0 and to switch 2 through a HBA in slot 0/1/3/0 on the system. The HP-UX OS in this configuration uses the native multi-pathing software “PV LINKS” to manage multiple paths to the same set of disks. For the configuration used in this document the same LUNs are available through both ports (a and c) on the storage array.




Figure 1: Topology Overview

Three different migration methods are discussed in the sections following the concepts section below. Each migration has its own pros and cons which are listed at the end of each section. The migration methods discussed in the document are:

Migrating Storage on HP-UX Servers To MDS Switches

1. Non Disruptive migration of storage and HP-UX host ports (the FC_ID (Domain_ID, Area_ID, and Port ID) changes during switch migration).
2. Non Disruptive migration of storage and HP-UX host ports (the FC_ID (Domain_ID, Area_ID, and Port ID) is retained during switch migration).
3. Disruptive migration of storage and HP-UX host ports across switches.

2 Concepts

In order to understand the various methods of migration discussed in this paper knowledge of certain Fibre Channel and HP-UX OS concepts is required. The required concepts are briefly described in this section (section 2).

2.1 Fibre Channel Identifier [FC_ID]

The Fibre Channel Identifier (FC_ID) is a 24 bit value that is assigned to every device that logs into the fabric and is unique for each device with in a fabric. The first 8 bits are reserved for the Domain_ID, next 8 bits for the Area_ID and the last 8 bits for the Port_ID as shown in Figure 2.




Figure 2: 24 Bit Fibre Channel Identifier

2.2 Volume Set addressing (VSA).

The parallel SCSI target allows 8 or 16 targets per controller as shown in Figure 3. In contrast, Fibre Channel has a huge address space and assigns 24 bit address to Fibre Channel nodes.




Figure 3: Parallel SCSI Device Addressing

Migrating Storage on HP-UX Servers To MDS Switches

With a 24 bit address a total of 2^{24} devices /entities can be addressed, which translates into $\sim 2^{24}$ SCSI targets (due to reserved address ranges, the actual number is 15,663,104). Furthermore the SCSI-3 protocol uses 8 bytes or 64 bits to address LUNs. Thus SCSI-3 can theoretically address $\sim 2^{24}$ targets each with $\sim 2^{64}$ LUNs through a single Fibre Channel controller.

To handle this huge address space HP-UX maps the Fibre Channel address space on to its Parallel SCSI address model. This model handles the large target address space of Fibre Channel by creating multiple virtual SCSI buses. On HP-UX each virtual bus addresses 16 targets and 8 LUNs per target. Thus each virtual bus can address 128 disks.

Volume Set Addressing (VSA) is one of the methods used by HP-UX to address a large number of LUNs per target and is typically used when addressing LUNs on large arrays. VSA uses a 14 bit number to address a volume (LUN). This addressing method can address a maximum of $2^{14} = 16384$ LUNs. As each virtual bus can address 128 LUNs, a total of 128 virtual buses are required with 14 bit addressing. Figure 4 shows the VSA on the array. It shows how many LUNs a target can potentially address using VSA.




Figure 4 : Volume set addressing on arrays

The 14 bit VSA address is split into 3 fields, the bus address, the target address and LUN address. The 14 bits are as show below.

- bits 13 – 7 bus address
- bits 6 – 3 target address
- bits 2 – 0 LUN address.

The Figure 5 shows the bit length of the various fields that make up VSA used by HP-UX OS.

Migrating Storage on HP-UX Servers To MDS Switches



Figure 5: 14 bit VSA address

As a random example, consider the Hardware path **1/8/0/0.120.0.22.1.15.0** that corresponds to the disk **cXt15d0**. The last 2 fields in the Hardware path define the target and LUN address. The volume address on the array for the device is derived from the last 3 fields of the Hardware path in this example 1.15.0. Here in this example the focus is on decoding the volume address (the device number on the array behind a Target). The value of “X” i.e. the controller number is automatically assigned by the kernel at the time of device creation.

Convert the last 3 fields in binary and into a 14 bit model defined above

1 15 0
0000001 1111 000 → convert to decimal value = 248 (LUN address on the array i.e. LUN number on the array behind Target 15)

E.g. 2: 0/0/4/0/0.103.0.11.0.11.5 → cXt11d5

0 11 5
0000000 1011 101 → decimal value = 93 (LUN address on the array, i.e. LUN number 95 behind Target 11 on the array.)

2.3 Logical Volume Manager (LVM)

This section briefly discusses Logical Volume Manager concepts that are used in this paper.

2.3.1 LVM basics

LVM is the volume management software on HP-UX systems. LVM uses Physical Volumes as its building blocks.

A disk is first initialized into a Physical Volume (PV). During initialization two reserved areas, PVRA (Physical Volume Reserved Area) and VGRA (Volume Group reserved Area) are created at the beginning of each physical volume (PV). These two reserved areas together are also called LVM header. This reserved space is used to store LVM (volume group and physical volume) related information. The size of these areas depends on the Volume Group and Logical Volume configuration parameters. It is usually > 3MB.

When the disk is initialized some space is also reserved for bad block pool. This space is used to map out bad blocks when they occur. The various reserved areas and the user data are on a physical disk is shown in Figure 6.

Migrating Storage on HP-UX Servers To MDS Switches




Figure 6: LVM Layout on a Physical Volume

The PVRA is unique for every PV in a Volume group. This contains the information about the

- Physical Volume: PV-ID, VG-ID, PE size etc.
- Start and the Length of: VGRA, user data and bad block pool
- MC Service Guard: Cluster Lock Area and Cluster ID
- Bad block directory for bad block pool

The VGRA on the other hand is same for all the PVs within a Volume Group. The VGRA, as the name suggests, contains information describing the Volume Group (VG). This area contains information about

- Volume Group: VG-ID, configured max_lv, max_pv, max_pe
- Per Logical Volume: LV flags, number mirrors, number of stripes, size etc.
- Per Physical Volume: PV-ID, PV size, PV flags, extent mappings etc.
- Volume group Status Area (VGSA): information regarding missing PVs and stale extents
- Mirror Consistency Record (MCR): For Mirror Write Cache handling

The VGRA is updated only at the time of volume group creation. It cannot be changed. If the values in the VGRA need to be changed the volume group has to be recreated.

PVRA is the area that needs to be updated when a CTD address of a PV changes to automatically import the Volume Group. The PVRA is updated only when `vgimport`, `vgextend` and `vgreduce` commands are executed against the PVs. This is why on a HP-UX system a combination of these commands is used to manipulate the Volume Group when the CTD address changes to facilitate auto import of Volume groups on system startup.

Migrating Storage on HP-UX Servers To MDS Switches

2.3.2 Physical Volume, Volume Group and Logical Volumes

A disk has to be initialized as a Physical Volume (PV) before being used under LVM. Once a PV is created it can then be added to a Volume Group. A PV can be member of only one Volume Group (VG) at a time.




Figure 7: Volume Group without PV Links

A volume group can consist of one or more PV(s). Each of the physical volumes can be addressed as cXtXdX (where c → Controller, t → Target and d → Disk (CTD)). Then Logical Volumes (LV) can be created from the PVs in the volume group. The LVs can be either be created on a single PV or striped across multiple PVs. When a LV is striped across multiple PV it mirror also needs to be striped across the same number of PVs.

Figure 7 shows the volume group and physical volume relationship with out PV links. Figure 8 shows the same with PV links. If the array has been configured for PV links only one of the devices can be initialized into a volume group. The alternate path can be added into the volume group using the **vgextend** command. Once the alternate path is added it shows up as an alternate link in the output of the **vgdisplay -v <VG Name>**.




Figure 8: Volume Group with PV Links

The Figure 9 shows the volume creation flow using HP-UX. A file system can be then created on these volumes and then mounted on the system for use.

Migrating Storage on HP-UX Servers To MDS Switches




Figure 9: LVM Volume Creation

Alternatively the volumes can be used as a raw volume as would be case with databases. A Sample vgdisplay is shown below.

```
sjhp2:/# vgdisplay -v vghpMDS
--- Volume groups ---
VG Name                /dev/vghpMDS
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 3
Open LV                 3
Max PV                 16
Cur PV                 5
Act PV                  5
Max PE per PV          2157
VGDA                   10
PE Size (Mbytes)       4
Total PE               10785
Alloc PE                7500
Free PE                 3285
Total PVG               0
Total Spare PVs         0
Total Spare PVs in use  0
```

Migrating Storage on HP-UX Servers To MDS Switches

--- Logical volumes ---

LV Name	/dev/vghpMDS/vol1
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol2
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol3
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1	
PV Name	/dev/dsk/c6t0d1	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d2	
PV Name	/dev/dsk/c6t0d2	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d3	
PV Name	/dev/dsk/c6t0d3	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d4	
PV Name	/dev/dsk/c6t0d4	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d5	
PV Name	/dev/dsk/c6t0d5	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	
sjhp2:/#		

2.3.3 LVM configuration file (/etc/lvmtab)

The /etc/lvmtab file contains information about all the known VGs and their related PVs. /etc/lvmtab is a binary data file. The viewable portions of the file can be viewed using the “**strings**” command. A sample output is as shown below. This file also contains (not viewable) VG-ID, total number of VGs, number of PVs per VG and status information. The lvmtab is automatically updated with the new VG and disk relationship information when VGs are created, removed or modified.

```
sjhp2# strings /etc/lvmtab
/dev/vg00
/dev/dsk/c2t0d0
/dev/vghpMDS
/dev/dsk/c4t0d1
/dev/dsk/c4t0d2
/dev/dsk/c4t0d3
/dev/dsk/c4t0d4
/dev/dsk/c4t0d5
/dev/dsk/c6t0d1
/dev/dsk/c6t0d2
/dev/dsk/c6t0d3
/dev/dsk/c6t0d4
/dev/dsk/c6t0d5
sjhp2#
```

All the VGs contained in the lvmtab file are automatically imported. This configuration file is used by the OS during start up to import the volume groups and activates them.

2.3.4 LVM limitations

LVM has the following limitations.

- The LVM configuration file (lvmtab) is a static file, i.e. it is updated only when a VG is created modified or deleted. CTD name changes (disk name change) are not automatically updated into the lvmtab.
- If any of the disks addresses (CTD) change, the import of the VG fails.
- LVM looks for the volume groups only on the PVs that are listed in /etc/lvmtab, in the above example it would look for VG “vghpMDS” only on the following disks: /dev/dsk/c4t0d1 /dev/dsk/c4t0d2 /dev/dsk/c4t0d3 /dev/dsk/c4t0d4 /dev/dsk/c4t0d5 /dev/dsk/c6t0d1 /dev/dsk/c6t0d2 /dev/dsk/c6t0d3 /dev/dsk/c6t0d4 /dev/dsk/c6t0d5.
- If the disk name (CTD) changes, the VG has to be imported manually.
- LVM has to be manually made aware of the PV link’s existence (using **vgextend**)

2.4 Multi path on HP-UX

Multi pathing on HP-UX is configured using PVLINKS. HP-UX uses active / passive multi pathing policy which allows only one path to be active at a given time, i.e. only one of the links can be used for I/O. The other path is strictly for failover. PVLINKS is the default multi path software on HP-UX. By default the OS fails over to the alternate path in 60 sec (if available). This parameter is configured in the kernel. Hence to change its value the kernel has to be modified and rebuilt. LVM cannot automatically detect the alternate path. PVLINKS have to be configured manually.

Migrating Storage on HP-UX Servers To MDS Switches

When a switch connected to HP-UX system is migrated / upgraded there is a chance that the disk address (CTD) may changes. When the disk address changes the LVM configurations has to be updated manually to reflect this change. Because of all the above limitations careful planning is needed for switch migration / upgrade when HP-UX hosts are involved.

2.5 Disk addressing, HP-UX and Fibre Channel HBAs.

HP-UX uses Volume Set Addressing (VSA) mode to address the large number of volumes (LUNs) on large arrays using Fibre Channel. HP-UX selects this mode of addressing of disks based on the SCSI **inquiry** data and LUN information returned by the SCSI-3 **REPORT LUN** command. This document assumes that VSA is being used on the storage array to enable a large number of LUNs to be visible through a target port on the storage array.

As explained above HP-UX creates multiple virtual SCSI buses in order to handle the large target address space of Fibre Channel. Each virtual bus can address up to 16 targets and each target can addresses 8 LUNs. Thus a virtual bus can address 128 LUNs. The Figure 10 shows device addressing on HP-UX server connected to large storage array.




Figure 10: Device addressing on HP-UX servers

The `ioscan` command is used to scan the bus for Hardware devices and also to list the devices on the system. A sample output of `ioscan` is as shown below.

```
sjhp2# ioscan -fnkC disk
Class      I  H/W Path          Driver    S/W State  H/W Type  Description
=====
disk       0  0/0/1/1.0.0       sdisk     CLAIMED    DEVICE    SEAGATE
ST318404LC /dev/dsk/c1t0d0    /dev/rdisk/c1t0d0
disk       2  0/0/2/0.0.0       sdisk     CLAIMED    DEVICE    SEAGATE
ST318404LC /dev/dsk/c2t0d0    /dev/rdisk/c2t0d0
disk       4  0/0/2/1.2.0       sdisk     CLAIMED    DEVICE    HP          DVD-ROM
304        /dev/dsk/c3t2d0    /dev/rdisk/c3t2d0
disk      12  0/1/0/0.111.0.2.0.0.0 sdisk     CLAIMED    DEVICE    EMC
SYMMETRIX  /dev/dsk/c5t0d0    /dev/rdisk/c5t0d0
```

Migrating Storage on HP-UX Servers To MDS Switches

```

disk      13  0/1/0/0.111.0.2.0.0.1  sdisk    CLAIMED    DEVICE    EMC
SYMMETRIX /dev/dsk/c5t0d1 /dev/rdisk/c5t0d1
disk      14  0/1/0/0.111.0.2.0.0.2  sdisk    CLAIMED    DEVICE    EMC
SYMMETRIX /dev/dsk/c5t0d2 /dev/rdisk/c5t0d2
disk      15  0/1/0/0.111.0.2.0.0.3  sdisk    CLAIMED    DEVICE    EMC
SYMMETRIX /dev/dsk/c5t0d3 /dev/rdisk/c5t0d3
  
```




Figure 11: Topology with Hardware paths

HP-UX device addressing is based on the Hardware path associated with the device, which means that every unique Hardware path to a device gets its own address in CxTx Dx format (C→ controller / channel, T→ target, D → disk/lun). The topology with the Hardware paths through two switches is shown in Figure 11.

The Hardware path as shown in Figure 12 is comprised of the HBA Hardware location, the FC_ID, the virtual bus, the target and the LUN. The Figure 12 shows one of the outputs from the ioscan command listing the contents of the output.

E.g.

Migrating Storage on HP-UX Servers To MDS Switches




Figure 12: ioscan -fnkC disk Output

Figure 11 shows the output of the `ioscan` command. The **man page** definitions of various fields of `ioscan` are as follows.

1. class → this includes disks, tapes, network cards etc
2. instance → it is a unique number assigned to a card or device within a class.
3. HW path → a numerical string of Hardware components, notated sequentially from the bus address to the device address.
4. driver → name of the driver that controls the device e.g. `sdisk` driver is used for disks, `stape` for tapes etc.
5. SW state → defines the status of the device. It can be `Claimed`, `NO_HW`, `error` etc.
6. HW Type → identifies the Hardware component. The types are `device` (disk), `memory`, `processor` etc.
7. Vendor → manufacturer of the Hardware
8. Block device → device file pointing to the block device (only on this device file system (block level) can be created and mounted.)
9. Character device → device file pointing to the character device. (device file can only be used to gain raw/character level access to the disk.)

In the sample `ioscan` output of the Hardware path shown in Figure 12 the HBA location is `0/0/4/0`, the `FC_ID` of the target port is `111.0.1` (in hex = `0x6F0001`) at bus 0, target 0, lun 0. The `CxTxDx` associated with this path is `c5t0d0`. In the above example when the `Domain_ID` is changed to 112 then the `FC_ID` changes to `112.0.1` (in hex = `0x700001`) and the Hardware path becomes `0/0/4/0.112.0.1.0.0.0` which is different from the example. This change forces the kernel to recognize it a new device. This path gets a new `CxTxDx`. In this example the disk address is `c7t0d0`. The Figure 13 shows the Hardware path before and after the `FC_ID` change.

As shown below in Figure 13 any change to the `FC_ID` of the target port results in the creation of a new Hardware path to the device. Even though the target and disk address of the device remains the same, the Hardware path changes with the `FC_ID`. To access any new Hardware path, a new device file has to be created to access it. Hence every change

Migrating Storage on HP-UX Servers To MDS Switches

in FC_ID associated with the target port in a fabric rennumbers the CTD address of devices on HP-UX servers.




Figure 13: Hardware Path's relationship to FC_ID on HP-UX

The controller number is assigned by the kernel at the time of device creation.

Target and LUN addresses are determined by the volume settings on the array, thus the “t” and “d” values are fixed. Thus in the above example the volume on the array would be the 0th volume. The target and disk address remains the same despite changes in the Hardware path to the device, and only the controller changes from c5 to c7 as see in Figure 14.

$$\text{Domain Id} + \text{Area} + \text{Port Number} = \text{FC_ID}$$




Figure 14: FCID and the Hardware path relation ship

FC_ID changes to the target port must be considered in the migration plans of a switch attached to a HP-UX system because the device (CxTx Dx) address changes with the FC_ID of the target port.

For reasons explained above the process of replacing a SAN switch attached to the HP-UX systems needs to be planned in detail.

3 Non Disruptive Migration – 1

It is always a good practice to backup all data before making any configuration changes to storage subsystems, SAN switches and the hosts.

3.1 Migrating Storage port independent of its FC_ID

During this migration process the CTD address of disks changes, but the migration of the target (storage ports on the array) and host to MDS switches is non disruptive.

At the beginning of the migration an HP-UX host is dual attached to the same set of disks through two Brocade switches, shown in Figure 15.

In phase one, one of the paths is disabled using LVM commands. The path through Brocade switch 2 becomes the active path as shown in Figure 16. Then the cable associated with Brocade switch 1 (disabled path) is disconnected and connected through MDS 1 as shown in Figure 17. The new devices are scanned and added as the alternate path to the path through the Brocade switch 2.

In phase two the path through the Brocade switch 2 is disabled. The path through MDS 1 takes over as the active path. The disabled path is then connected through MDS 2 and the new devices scanned and added as alternate path(s) to the path(s) through MDS 1. Both paths are now through MDS 1 and MDS 2 as shown in Figure 18. Figure 15 shows the topology at the start of the migration.




Figure 15: Initial Configuration

The Domain_IDs used in the Brocade switches are 4 and 5 for switch 1 and switch 2 respectively. The storage ports are connected to port 0 on both switches. The FC_ID of the targets as assigned by the switches are 0x040000 on switch 1 and 0x050000 on switch 2.

Migrating Storage on HP-UX Servers To MDS Switches

3.1.1 Migration Steps

The non disruptive migration of storage attached to a HP-UX server through 3rd party switches to MDS switches must be done in 2 phases, each of which has 3 steps.

➤ Step 1

- Create an inventory of the storage in the system
- Identify the alternate paths, identify the relationship between disks in a volume group, identifies the relationship between disk and volume groups
- Finally remove one of the alternate paths. The I/O on the system temporarily goes through one path.

➤ Step 2

- Physically remove fiber cable from the disabled alternate path
- Plug the disconnected fiber cables in to the MDS switch.

➤ Step 3

- Create the appropriate VSAN and zone configurations on the switch.
- Run **ioscan -fnC disk** to discover the devices on the path through the MDS switch.
- Run **insf -eC disk** to create the new device files.
- Add the alternate path and the system has path redundancy again.

*Note: LVM only allows one path to be active at a time. Hence to make an inactive path active, simply deactivate the active path (using **vgreduce**). This makes the inactive path active. Similarly when a second path is added to the existing path it is added as an inactive path only. No I/O is lost as LVM takes care of rerouting I/O through the other available path which is now active.*

➤ Step 4

Repeat steps 1, 2, 3 on the other path connected through the Brocade switch. Once this is done the both switches are MDS switches. This migration can be done online with out any down time for the application.

3.1.2 Step 1: Deactivate one of the redundant paths

Current configuration is as shown in Figure 15. The storage is currently visible through 2 Brocade switches.

In this step, information regarding the various devices and their relationships (alternate path) to each other and the volume groups is collected. One of the redundant paths is disabled. There are two volume groups on sjhp2, vg00 and vghpMDS. vg00 is on a local disk and is used to manage the root disk only. vghpMDS consists of 6 disks. All seven disks have an alternate path. The **ioscan** command is used to identify the relationship

Migrating Storage on HP-UX Servers To MDS Switches

between disks and volume groups. In step 1, the commands used are **ioscan**, **vgdisplay** and **vgreduce**.

3.1.2.1 Identify the Hardware path involved for migration

Run **ioscan -fnkC disk** and redirect it to a file. This output is used as a reference for the entire migration. **ioscan** queries the kernel for disks and lists them based on their Hardware path. In the output below the Hardware paths have been highlighted. The command **ioscan -fnkC disk** lists three Hardware paths /0/0/2, 0/1/1, and 0/1/3 with devices associated with them. 0/0/2 is the internal SCSI controller while 0/1/1 and 0/1/3 are FC HBAs.

```
sjhp2:/# ioscan -fnkC disk
Class      I  H/W Path      Driver    S/W State  H/W Type  Description
=====
disk       0  0/0/2/0.0.0    sdisk     CLAIMED    DEVICE     SEAGATE
ST318404LC /dev/dsk/c2t0d0 /dev/rdisk/c2t0d0
disk       1  0/0/2/1.2.0    sdisk     CLAIMED    DEVICE     HP          DVD-ROM
304        /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk       2  0/1/1/0.4.0.0.0.0.0 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d0 /dev/rdisk/c4t0d0
disk       3  0/1/1/0.4.0.0.0.0.1 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d1 /dev/rdisk/c4t0d1
disk       4  0/1/1/0.4.0.0.0.0.2 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d2 /dev/rdisk/c4t0d2
disk       5  0/1/1/0.4.0.0.0.0.3 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d3 /dev/rdisk/c4t0d3
disk       6  0/1/1/0.4.0.0.0.0.4 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d4 /dev/rdisk/c4t0d4
disk       7  0/1/1/0.4.0.0.0.0.5 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d5 /dev/rdisk/c4t0d5
disk       8  0/1/3/0.5.0.0.0.0.0 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d0 /dev/rdisk/c6t0d0
disk       9  0/1/3/0.5.0.0.0.0.1 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d1 /dev/rdisk/c6t0d1
disk      10  0/1/3/0.5.0.0.0.0.2 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d2 /dev/rdisk/c6t0d2
disk      11  0/1/3/0.5.0.0.0.0.3 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d3 /dev/rdisk/c6t0d3
disk      12  0/1/3/0.5.0.0.0.0.4 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d4 /dev/rdisk/c6t0d4
disk      13  0/1/3/0.5.0.0.0.0.5 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d5 /dev/rdisk/c6t0d5
```

From the **ioscan -fnkC disk** output the Hardware path 0/1/1/0.4.0.0.0.0.0, the HBA located at 0/1/1/0 is connected to Brocade switch 1 with Domain_ID on the switch set to 4. Similarly the Hardware path 0/1/3/0.5.0.0.0.0.5, the HBA located at 0/1/3/0 is connected to Brocade switch 2 with Domain_ID on the switch set to 5.

The command **ioscan -fnkC fc** lists all the fibre channel controllers on the system and their Hardware paths.

Migrating Storage on HP-UX Servers To MDS Switches

```
sjhp2:/# ioscan -fnkC fc
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
fc	0	0/1/1/0	td	CLAIMED	INTERFACE	HP Tachyon XL2 Fibre Channel Mass Storage Adapter /dev/td0
fc	1	0/1/2/0	td	CLAIMED	INTERFACE	HP Tachyon XL2 Fibre Channel Mass Storage Adapter dev/td1
fc	2	0/1/3/0	td	CLAIMED	INTERFACE	HP Tachyon XL2 Fibre Channel Mass Storage Adapter /dev/td2

The command **ioscan -fnkC fc** shows that three FC HBAs on sjhp2. They are on Hardware paths 0/1/1, 0/1/2 and 0/1/3. **ioscan -fnkC disk** show disks only on two of the 3 paths, on 0/1/1 and 0/1/3. The disks c4t0d0 d5 are visible to the HBA on Hardware path 0/1/1 and devices c6t0d1-d5 are visible to the HBA on Hardware path 0/1/3. The HBA on Hardware path 0/1/2 does not see any disk on sjhp2.

vgdisplay -v lists details about all volume groups on a particular host. The command **vgdisplay -v vghpMDS** lists only the details about the volume group vghpMDS.

```
sjhp2:/# vgdisplay -v vghpMDS
```

```
--- Volume groups ---
VG Name                /dev/vghpMDS
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 3
Open LV                 3
Max PV                  16
Cur PV                  5
Act PV                  5
Max PE per PV           2157
VGDA                     10
PE Size (Mbytes)        4
Total PE                10785
Alloc PE                 7500
Free PE                  3285
Total PVG                0
Total Spare PVs          0
Total Spare PVs in use  0

--- Logical volumes ---
LV Name                /dev/vghpMDS/vol1
LV Status               available/syncd
LV Size (Mbytes)        10000
Current LE              2500
Allocated PE             2500
Used PV                  5

LV Name                /dev/vghpMDS/vol2
LV Status               available/syncd
LV Size (Mbytes)        10000
Current LE              2500
Allocated PE             2500
Used PV                  5

LV Name                /dev/vghpMDS/vol3
LV Status               available/syncd
LV Size (Mbytes)        10000
Current LE              2500
Allocated PE             2500
```

Migrating Storage on HP-UX Servers To MDS Switches

Used PV

5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1	
PV Name	/dev/dsk/c6t0d1	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d2	
PV Name	/dev/dsk/c6t0d2	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d3	
PV Name	/dev/dsk/c6t0d3	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d4	
PV Name	/dev/dsk/c6t0d4	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d5	
PV Name	/dev/dsk/c6t0d5	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

sjhp2:/#

The output has three sections, Volume group, Logical volumes and Physical volumes. The Physical volumes section provides information regarding the disk and alternate paths if any are configured. This information is required for the migration. The above output shows the PVLINK setup for the disks. The devices c4t0d1-d5 are on the primary path and the devices c6t0d1-d5 are on the alternate path. The alternate path needs to be removed in preparation for migration over to the MDS.

Migrating Storage on HP-UX Servers To MDS Switches

3.1.2.2 Remove the redundant path




Figure 16: Remove the path through Brocade

The following command, **vgreduce /dev/vghpMDS /dev/dsk/c6t0d1 /dev/dsk/c6t0d2 /dev/dsk/c6t0d3 /dev/dsk/c6t0d4 /dev/dsk/c6t0d5** removes the listed disks from the volume group vghpMDS. The disks are still visible to the sjhp2. The redundant path is removed for the disks in the volume group vghpMDS. Disconnect the cables from switch 2. Figure 16 shows the scenario when the cables are disconnected from switch 2.

```
sjhp2:/# vgreduce /dev/vghpMDS /dev/dsk/c6t0d1 /dev/dsk/c6t0d2 /dev/dsk/c6t0d3
/dev/dsk/c6t0d4 /dev/dsk/c6t0d5
Device file path "/dev/dsk/c6t0d1" is an alternate path.
Device file path "/dev/dsk/c6t0d2" is an alternate path.
Device file path "/dev/dsk/c6t0d3" is an alternate path.
Device file path "/dev/dsk/c6t0d4" is an alternate path.
Device file path "/dev/dsk/c6t0d5" is an alternate path.
Volume group "/dev/vghpMDS" has been successfully reduced.
Volume Group configuration for /dev/vghpMDS has been saved in
/etc/lvmconf/vghpMDS.conf
sjhp2:/#
```

Verify that the alternate link has been removed. **vgdisplay** of the volume group vghpMDS lists the current status of the disks in the volume group. The output below shows only a single path to the disks in the volume group vghpMDS.

```
sjhp2:/# vgdisplay -v /dev/vghpMDS
sjhp2:/dev/vghpMDS# vgdisplay -v vghpMDS
--- Volume groups ---
VG Name                /dev/vghpMDS
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 3
Open LV                 3
Max PV                 16
```

Migrating Storage on HP-UX Servers To MDS Switches

Cur PV	5
Act PV	5
Max PE per PV	2157
VGDA	10
PE Size (Mbytes)	4
Total PE	10785
Alloc PE	7500
Free PE	3285
Total PVG	0
Total Spare PVs	0
Total Spare PVs in use	0

--- Logical volumes ---

LV Name	/dev/vghpMDS/vol1
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol2
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol3
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d2
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d3
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d4
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d5
PV Status	available
Total PE	2157

Migrating Storage on HP-UX Servers To MDS Switches

```
Free PE
Autoswitch
sjhp2:/#
```

```
657
On
```

The Physical volumes section no longer shows the alternate link as being present. The path associated with the c6 disks has now been removed. The system is now running on only one path, the path associated with disks c4. Now the cables can be disconnected and moved over to the MDS.

3.1.3 Step 2: Move the disconnected path MDS

The new configuration is shown in Figure 17. One path to the storage is through a MDS switch and the other path is through a Brocade switch.




Figure 17: MDS Switch in the Fabric

In this step the cables associated with the removed Hardware path 0/1/3 are removed from the Brocade switch and migrated over to the MDS switch. The MDS switch is configured to allow the host to see the storage through it. Then the IO subsystem on sjhp2 is scanned for the new devices. In step 2 the commands used are **ioscan** and **vgdisplay** on the HP-UX hosts.

3.1.4 Step 3: Scan and list the current status of the devices

The **ioscan -fnC disk** command is run to rescan the IO bus and look for new devices on the host after the path through the Brocade is disconnected. The output of the **ioscan** command show that the devices associated with the Hardware path 0/1/3 are no longer visible to the host. The lines with NO_HW are the disks that are no longer visible. The Hardware path and the device status are highlighted in red.

```
sjhp2:/# ioscan -fnC disk
Class      I  H/W Path      Driver  S/W State  H/W Type  Description
=====
disk       0  0/0/2/0.0.0    sdisk   CLAIMED    DEVICE    SEAGATE
ST318404LC /dev/dsk/c2t0d0 /dev/rdisk/c2t0d0
disk       1  0/0/2/1.2.0    sdisk   CLAIMED    DEVICE    HP        DVD-ROM
304        /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
```

Migrating Storage on HP-UX Servers To MDS Switches

disk	2	0/1/1/0.4.0.0.0.0.0	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d0	/dev/rdisk/c4t0d0			
disk	3	0/1/1/0.4.0.0.0.0.1	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d1	/dev/rdisk/c4t0d1			
disk	4	0/1/1/0.4.0.0.0.0.2	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d2	/dev/rdisk/c4t0d2			
disk	5	0/1/1/0.4.0.0.0.0.3	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d3	/dev/rdisk/c4t0d3			
disk	6	0/1/1/0.4.0.0.0.0.4	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d4	/dev/rdisk/c4t0d4			
disk	7	0/1/1/0.4.0.4.0.0.5	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d5	/dev/rdisk/c4t0d5			
disk	8	0/1/3/0.5.0.0.0.0.0	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d0	/dev/rdisk/c6t0d0			
disk	9	0/1/3/0.5.0.0.0.0.1	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d1	/dev/rdisk/c6t0d1			
disk	10	0/1/3/0.5.0.0.0.0.2	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d2	/dev/rdisk/c6t0d2			
disk	11	0/1/3/0.5.0.0.0.0.3	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d3	/dev/rdisk/c6t0d3			
disk	12	0/1/3/0.5.0.0.0.0.4	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d4	/dev/rdisk/c6t0d4			
disk	13	0/1/3/0.5.0.0.0.0.5	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d5	/dev/rdisk/c6t0d5			

sjhp2:/#

3.1.4.1 MDS switch setup

Connect the HBA and the storage port to MDS1. On the MDS switch create a VSAN; configure the Domain_ID to be static; configure the FC_ID to be persistent. This is a very important step for the HP-UX hosts. This prevents the disks on the HP from getting new CTD numbers if the switch reboots. HP-UX assigns the CTD numbers for a disks based on the Hardware location of the HBA, FC_ID and Volume address of the lun on the array. Zone the storage (target) port and the HBA. The host sees the same set of LUNS, once through MDS1 and the second time through the Brocade as shown in Figure 17.

3.1.4.2 Create device files for the new devices in order to access them.

ioscan -fnc disk now scans the I/O bus and rediscovers the devices that are now visible through MDS1. Since the Domain_ID and the FC_ID for the storage port is different on MDS1 (than on the Brocade) the host sees the same disk with new CTD numbers. The Hardware location 0/1/3 of the HBA still remains the same. The **insf -ec disk** creates the special files needed to make use of the disks accessible to the application (LVM) on sjhp2.

In the output below the disks in **green** are the devices seen through the Brocade. The entries in **red** are the entries for the disks that were on the disconnected (Brocade) path. The entries in **blue** are the disks (*same as the ones seen through Brocade*) seen through MDS1.

Migrating Storage on HP-UX Servers To MDS Switches

```
sjhp2:/# ioscan -fnC disk
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	0	0/0/2/0.0.0	sdisk	CLAIMED	DEVICE	SEAGATE
ST318404LC		/dev/dsk/c2t0d0	/dev/rdisk/c2t0d0			
disk	1	0/0/2/1.2.0	sdisk	CLAIMED	DEVICE	HP DVD-ROM
304		/dev/dsk/c3t2d0	/dev/rdisk/c3t2d0			
disk	2	0/1/1/0.4.0.0.0.0	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d0	/dev/rdisk/c4t0d0			
disk	3	0/1/1/0.4.0.0.0.1	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d1	/dev/rdisk/c4t0d1			
disk	4	0/1/1/0.4.0.0.0.2	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d2	/dev/rdisk/c4t0d2			
disk	5	0/1/1/0.4.0.0.0.3	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d3	/dev/rdisk/c4t0d3			
disk	6	0/1/1/0.4.0.0.0.4	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d4	/dev/rdisk/c4t0d4			
disk	7	0/1/1/0.4.0.0.0.5	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c4t0d5	/dev/rdisk/c4t0d5			
disk	8	0/1/3/0.5.0.0.0.0	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d0	/dev/rdisk/c6t0d0			
disk	9	0/1/3/0.5.0.0.0.1	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d1	/dev/rdisk/c6t0d1			
disk	10	0/1/3/0.5.0.0.0.2	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d2	/dev/rdisk/c6t0d2			
disk	11	0/1/3/0.5.0.0.0.3	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d3	/dev/rdisk/c6t0d3			
disk	12	0/1/3/0.5.0.0.0.4	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d4	/dev/rdisk/c6t0d4			
disk	13	0/1/3/0.5.0.0.0.5	sdisk	NO_HW	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d5	/dev/rdisk/c6t0d5			
disk	14	0/1/3/0.111.0.0.0.0	sdisk	CLAIMED	DEVICE	EMC SYMMETRIX
disk	15	0/1/3/0.111.0.0.0.1	sdisk	CLAIMED	DEVICE	EMC SYMMETRIX
disk	16	0/1/3/0.111.0.0.0.2	sdisk	CLAIMED	DEVICE	EMC SYMMETRIX
disk	17	0/1/3/0.111.0.0.0.3	sdisk	CLAIMED	DEVICE	EMC SYMMETRIX
disk	18	0/1/3/0.111.0.0.0.4	sdisk	CLAIMED	DEVICE	EMC SYMMETRIX
disk	19	0/1/3/0.111.0.0.0.5	sdisk	CLAIMED	DEVICE	EMC SYMMETRIX

```
sjhp2:/#
```

```
sjhp2:/# insf -eC disk
```

```
insf: Installing special files for sdisk instance 0 address 0/0/2/0.0.0
insf: Installing special files for sdisk instance 1 address 0/0/2/1.2.0
insf: Installing special files for sdisk instance 2 address 0/1/1/0.4.0.0.0.0
insf: Installing special files for sdisk instance 3 address 0/1/1/0.4.0.0.0.1
insf: Installing special files for sdisk instance 4 address 0/1/1/0.4.0.0.0.2
insf: Installing special files for sdisk instance 5 address 0/1/1/0.4.0.0.0.3
insf: Installing special files for sdisk instance 6 address 0/1/1/0.4.0.0.0.4
insf: Installing special files for sdisk instance 7 address 0/1/1/0.4.0.0.0.5
insf: Installing special files for sdisk instance 8 address 0/1/3/0.5.0.0.0.0
insf: Installing special files for sdisk instance 9 address 0/1/3/0.5.0.0.0.1
insf: Installing special files for sdisk instance 10 address
0/1/3/0.5.0.0.0.2
insf: Installing special files for sdisk instance 11 address
0/1/3/0.5.0.0.0.3
insf: Installing special files for sdisk instance 12 address
0/1/3/0.5.0.0.0.4
insf: Installing special files for sdisk instance 13 address
0/1/3/0.5.0.0.0.5
insf: Installing special files for sdisk instance 14 address
0/1/3/0.111.0.0.0.0
insf: Installing special files for sdisk instance 15 address
0/1/3/0.111.0.0.0.1
```

Migrating Storage on HP-UX Servers To MDS Switches

```
insf: Installing special files for sdisk instance 16 address
0/1/3/0.111.0.0.0.0.2
insf: Installing special files for sdisk instance 17 address
0/1/3/0.111.0.0.0.0.3
insf: Installing special files for sdisk instance 18 address
0/1/3/0.111.0.0.0.0.4
insf: Installing special files for sdisk instance 19 address
0/1/3/0.111.0.0.0.0.5
sjhp2:/#
```

The **ioscan -fnkC disk** shows the disks seen through the Brocade path in **green**, the disconnected devices (these go away when the system is rebooted or when the **rmsf** command is run to remove the defunct disks) are highlighted **red** and the new devices seen through the MDS switch are highlighted in **blue**.

```
sjhp2:/# ioscan -fnkC disk
Class      I  H/W Path      Driver      S/W State   H/W Type     Description
=====
disk       0  0/0/2/0.0.0    sdisk       CLAIMED      DEVICE       SEAGATE
ST318404LC /dev/dsk/c2t0d0 /dev/rdisk/c2t0d0
disk       1  0/0/2/1.2.0    sdisk       CLAIMED      DEVICE       HP           DVD-ROM
304        /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk      245  0/1/1/0.4.0.0.0.0.0 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c4t0d0 /dev/rdisk/c4t0d0
disk      246  0/1/1/0.4.0.0.0.0.1 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c4t0d1 /dev/rdisk/c4t0d1
disk      247  0/1/1/0.4.0.0.0.0.2 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c4t0d2 /dev/rdisk/c4t0d2
disk      248  0/1/1/0.4.0.0.0.0.3 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c4t0d3 /dev/rdisk/c4t0d3
disk      249  0/1/1/0.4.0.0.0.0.4 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c4t0d4 /dev/rdisk/c4t0d4
disk      250  0/1/1/0.4.0.0.0.0.5 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c4t0d5 /dev/rdisk/c4t0d5
disk      292  0/1/3/0.5.0.0.0.0.0 sdisk       NO_HW        DEVICE       EMC
SYMMETRIX  /dev/dsk/c6t0d0 /dev/rdisk/c6t0d0
disk      293  0/1/3/0.5.0.0.0.0.1 sdisk       NO_HW        DEVICE       EMC
SYMMETRIX  /dev/dsk/c6t0d1 /dev/rdisk/c6t0d1
disk      294  0/1/3/0.5.0.0.0.0.2 sdisk       NO_HW        DEVICE       EMC
SYMMETRIX  /dev/dsk/c6t0d2 /dev/rdisk/c6t0d2
disk      295  0/1/3/0.5.0.0.0.0.3 sdisk       NO_HW        DEVICE       EMC
SYMMETRIX  /dev/dsk/c6t0d3 /dev/rdisk/c6t0d3
disk      296  0/1/3/0.5.0.0.0.0.4 sdisk       NO_HW        DEVICE       EMC
SYMMETRIX  /dev/dsk/c6t0d4 /dev/rdisk/c6t0d4
disk      297  0/1/3/0.5.0.0.0.0.5 sdisk       NO_HW        DEVICE       EMC
SYMMETRIX  /dev/dsk/c6t0d5 /dev/rdisk/c6t0d5
disk      339  0/1/3/0.111.0.0.0.0.0 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c8t0d0 /dev/rdisk/c8t0d0
disk      340  0/1/3/0.111.0.0.0.0.1 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c8t0d1 /dev/rdisk/c8t0d1
disk      341  0/1/3/0.111.0.0.0.0.2 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c8t0d2 /dev/rdisk/c8t0d2
disk      342  0/1/3/0.111.0.0.0.0.3 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c8t0d3 /dev/rdisk/c8t0d3
disk      343  0/1/3/0.111.0.0.0.0.4 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c8t0d4 /dev/rdisk/c8t0d4
disk      344  0/1/3/0.111.0.0.0.0.5 sdisk       CLAIMED      DEVICE       EMC
SYMMETRIX  /dev/dsk/c8t0d5 /dev/rdisk/c8t0d5
```

Migrating Storage on HP-UX Servers To MDS Switches

3.1.4.3 Add alternate path to Volume group.

The OS identifies the alternate path based on unique serial number assigned to each disk. In the case of large arrays the array assigns a unique serial number to each volume/LUN that is exposed to a HBA. If two Hardware paths point to a disk with the same serial number then they are considered to be alternate paths to each other. The HP-UX has the capability of identifying the alternate paths using the serial number of the disks. Thus when a disk is added to a volume group, the OS determines whether it is an alternate path. If it is an alternate path, the device is added as an alternate path. Otherwise it is added as additional storage to the disk group.

ioscan -fnkC disk shows six new devices c8t0d1-d5. The output shows that the new disks are attached to the HBA on Hardware location 0/1/3. These are the very same devices that were visible as c6t0d1-d5 when the HBA was connected through the Brocade switch. Sections of the Hardware path have changed. The old path was 0/1/3/0.5.0.0.0.0 (the highlighted positions corresponding to the FC_ID of the target port.) and the new Hardware path is 0/1/3/0.111.0.0.0.1. Since the Hardware path has changed the OS registers it a new device and creates a new device file for it. HP-UX derives the CTD numbers for the disks based on the FC_ID of the storage port. The FC_ID of the storage port changed when the storage port was moved over to the switch MDS1. The Domain_ID was 5 in the Brocade switch 1 and is 111 on the switch MDS1. This is reflected in the Hardware path of the disk. Using **vgextend** disks c8t0d1-d5 are added to the volume group vghpMDS. In this instance SYMCLI commands are used to identify the devices with the same serial numbers.

The disks c8t0d1-d5 are alternate path to the c4 disks, they are added as alternate links. These are the disks that were visible under controller c6. This is further confirmed since the devices are associated with the HBA located at 0/1/3/0 on the server.

The command **vgextend /dev/vghpMDS /dev/dsk/c8t0d1 /dev/dsk/c8t0d2 /dev/dsk/c8t0d3 /dev/dsk/c8t0d4 /dev/dsk/c8t0d5** adds the disks as alternate links.

```
sjhp2:/# vgextend /dev/vghpMDS /dev/dsk/c8t0d1 /dev/dsk/c8t0d2 /dev/dsk/c8t0d3
/dev/dsk/c8t0d4 /dev/dsk/c8t0d5
Volume group "/dev/vghpMDS" has been successfully extended.
Volume Group configuration for /dev/vghpMDS has been saved in
/etc/lvmconf/vghpMDS.conf
sjhp2:/#
```

Run **vgdisplay** command to verify that the alternate path is now in place through the MDS switch.

```
sjhp2:/# vgdisplay -v vghpMDS
--- Volume groups ---
VG Name                /dev/vghpMDS
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 3
Open LV                 3
```

Migrating Storage on HP-UX Servers To MDS Switches

Max PV	16
Cur PV	5
Act PV	5
Max PE per PV	2157
VGDA	10
PE Size (Mbytes)	4
Total PE	10785
Alloc PE	7500
Free PE	3285
Total PVG	0
Total Spare PVs	0
Total Spare PVs in use	0

--- Logical volumes ---

LV Name	/dev/vghpMDS/vol1
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol2
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol3
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1
PV Name	/dev/dsk/c8t0d1 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d2
PV Name	/dev/dsk/c8t0d2 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d3
PV Name	/dev/dsk/c8t0d3 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d4
PV Name	/dev/dsk/c8t0d4 Alternate Link
PV Status	available
Total PE	2157
Free PE	657

Migrating Storage on HP-UX Servers To MDS Switches

```
Autoswitch                On

PV Name                    /dev/dsk/c4t0d5
PV Name                    /dev/dsk/c8t0d5 Alternate Link
PV Status                  available
Total PE                   2157
Free PE                    657
Autoswitch                 On

sjhp2:/#
```

From the above output it can be seen that the Alternate links which were through the Brocade switch i.e. devices c6t0d1-d5 are now pointing to devices c8t0d1 ... d6 which are visible through MDS1. The alternate link is highlighted blue. One path has been successfully migrated to the MDS switch.

3.1.5 Step 4: Replace the other Brocade to a full MDS fabric.




Figure 18: All MDS Topology

Repeat all the above process that was used to replace Brocade switch 1 to replace Brocade switch 2. Use `vgreduce` to remove the disks c2t0d1-d5 from the volume group `vgMDS`. Then disconnect the path through Brocade switch 2. This path is then reconnected through MDS2. When the disks c4t0d1-d5 are removed from the volume group `vgMDS` the path fails over to c8t0d1-d5 through MDS1 with out interruption to the services. Then run `ioscan -fnc disk` to scan for new devices that will be visible through the switch MDS2. Then use `vgextend` to add the new disks as alternate path to the devices seen through the switch MDS1. This completes the migration of the HP-UX host off of the Brocade switches non-disruptively. The final topology after migration is shown in Figure 18.

3.2 Pros and cons

3.2.1 Pros

- The whole migration is non disruptive and simple
- This whole process can be done online (while the host and applications are still up and running)
- Only one path needs to be migrated at a time. This leaves the alternate path always available for performing I/O to the system.
- Fail back to the old infrastructure is relatively simple.

3.2.2 Cons

- The CTD addresses of the disks changes.
- The old device files (the ones with the NO_HW tag) entries are still visible on the host. (`rmsf -a -v <Hardware path>`) or a reboot at a later time is required to remove defunct device files.
- This migration cannot be performed on a cluster (MC Service Guard - HP cluster software) running Oracle parallel server. This is because the cluster has to be taken down is the CTD addresses need to change.

4 Non Disruptive Migration – 2 (using the persistent FC_ID feature)

4.1 FC_ID remains persistent across the switches (before and after migration).

In this migration example, the move is transparent. The CTD addresses of the disks remain the same before and after the migration to MDS switches.

In this topology one of the paths to the disk on the SJHP2 is through a McData switch and the other path is through a Brocade switch. This topology is selected to show the Domain_ID numbering differences between Brocade and McData and how it affects the migration to MDS. A HP-UX host is dual attached to the same set of disks through a Brocade switch and a McData switch. In phase one, the path through the Brocade is disabled using LVM commands. The path through McData becomes the active path. Then the cable associated with the Brocade (disabled path) is disconnected and reconnected through MDS 1. The new devices are scanned and added as an alternate path to the path through the McData switch. In phase two the path through the McData switch is disabled. The path through MDS 1 takes over as the active path. The disabled path is then reconnected through MDS 2 and the new devices scanned and added as an alternate path to the path through MDS 1. Both paths are now connected through MDS switches 1 and 2.

The Domain_ID of the Brocade is 4 and the storage is connected to port 0 and the host is connected to port 1. The Domain_ID of the McData is 1 and the storage is connected to port 0 and the host is connected to port 1. The Hardware path for disks attached to HBA in slot 0/1/1/0 connected through McData is 0/1/1/0.97.4.19. The Hardware path for disks attached to HBA in slot 0/1/3/0 connected through Brocade is 0/1/3/0.4.0.0 This configuration is shown in Figure 19.




Figure 19: Topology with Brocade and McData switches

Migrating Storage on HP-UX Servers To MDS Switches

4.2 Migrating the ports on Brocade to MDS

4.2.1 Gather data before removing the path through Brocade.

Before this migration can begin, information regarding the various disk groups is gathered. The **ioscan** and **vgdisplay** commands are used to gather the required information. Figure 19 shows the initial topology before the migration.

```
sjhp2:/# ioscan -fnkc disk
Class      I  H/W Path      Driver    S/W State  H/W Type    Description
=====
disk       0  0/0/2/0.0.0    sdisk     CLAIMED    DEVICE      SEAGATE
ST318404LC /dev/dsk/c2t0d0 /dev/rdisk/c2t0d0
disk       1  0/0/2/1.2.0    sdisk     CLAIMED    DEVICE      HP          DVD-ROM
304        /dev/dsk/c3t2d0 /dev/rdisk/c3t2d0
disk       2  0/1/1/0.97.4.19.0.0.0 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c4t0d0 /dev/rdisk/c4t0d0
disk       3  0/1/1/0.97.4.19.0.0.1 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c4t0d1 /dev/rdisk/c4t0d1
disk       4  0/1/1/0.97.4.19.0.0.2 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c4t0d2 /dev/rdisk/c4t0d2
disk       5  0/1/1/0.97.4.19.0.0.3 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c4t0d3 /dev/rdisk/c4t0d3
disk       6  0/1/1/0.97.4.19.0.0.4 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c4t0d4 /dev/rdisk/c4t0d4
disk       7  0/1/1/0.97.4.19.0.0.5 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c4t0d5 /dev/rdisk/c4t0d5
disk       8  0/1/3/0.4.0.0.0.0.0 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c6t0d0 /dev/rdisk/c6t0d0
disk       9  0/1/3/0.4.0.0.0.0.1 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c6t0d1 /dev/rdisk/c6t0d1
disk      10  0/1/3/0.4.0.0.0.0.2 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c6t0d2 /dev/rdisk/c6t0d2
disk      11  0/1/3/0.4.0.0.0.0.3 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c6t0d3 /dev/rdisk/c6t0d3
disk      12  0/1/3/0.4.0.0.0.0.4 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c6t0d4 /dev/rdisk/c6t0d4
disk      13  0/1/3/0.4.0.0.0.0.5 sdisk     CLAIMED    DEVICE      EMC
SYMMETRIX  /dev/dsk/c6t0d5 /dev/rdisk/c6t0d5
```

From the **ioscan** output it can be seen that, based on the FC_ID portion (highlighted in blue) of the Hardware path, the Hardware path for disks c4t0d0 ... d7 is the path through the McData switch (Domain_ID 1). Similarly the Hardware path for disks c6t0d0 ... d7 is the path through the Brocade switch (Domain_ID 4).

```
sjhp2:/dev/vghpMDS# vgdisplay -v vghpMDS
--- Volume groups ---
VG Name                /dev/vghpMDS
VG Write Access        read/write
VG Status              available
Max LV                 255
Cur LV                3
Open LV                3
Max PV                 16
Cur PV                5
Act PV                 5
Max PE per PV          2157
VGDA                   10
PE Size (Mbytes)       4
```

Migrating Storage on HP-UX Servers To MDS Switches

Total PE	10785
Alloc PE	7500
Free PE	3285
Total PVG	0
Total Spare PVs	0
Total Spare PVs in use	0

--- Logical volumes ---

LV Name	/dev/vghpMDS/lvol1
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/lvol2
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/lvol3
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1
PV Name	/dev/dsk/c6t0d1 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d2
PV Name	/dev/dsk/c6t0d2 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d3
PV Name	/dev/dsk/c6t0d3 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d4
PV Name	/dev/dsk/c6t0d4 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d5
PV Name	/dev/dsk/c6t0d5 Alternate Link
PV Status	available
Total PE	2157

Migrating Storage on HP-UX Servers To MDS Switches

Free PE
Autoswitch

657
On

The above `vgdisplay` output shows that the disks `c4t0d0 ... d7` belong to the primary path which is active and `c6t0d0 ... d7` belong to the alternate path which is passive. In the next step the disks `c6t0d1-d5` are removed from the disk group `vggrpMDS` in preparation for the migration to MDS switch.




Figure 20: Path through Brocade is disconnected.

The path through the Brocade switch needs to be migrated over to MDS1. Before migration of the path the FC_ID and the Domain_ID information from the Brocade switch is needed, this can be gleaned with the `switchShow` and `nsShow` commands.

```
Brocade:admin> switchShow
switchName:      Bot3800
switchType:      9.2
switchState:     Online
switchMode:      Native
switchRole:      Subordinate
switchDomain:     4          ← Domain_ID of  Brocade
switchId:        fffc04
switchHardwarewn: 10:00:00:60:69:51:3a:60
switchBeacon:    OFF
Zoning:          ON (s1)

port 0: id N2 Online      F-Port 50:06:04:82:bf:d1:db:cd
port 1: id N2 Online      F-Port 50:06:0b:00:00:11:b9:18
port 2: -- N2 No_Module
port 3: -- N2 No_Module
port 4: -- N2 No_Module
port 5: -- N2 No_Module
port 6: -- N2 No_Module
port 7: -- N2 No_Module
port 8: -- N2 No_Module
port 9: -- N2 No_Module
port 10: -- N2 No_Module
port 11: -- N2 No_Module
port 12: -- N2 No_Module
port 13: -- N2 No_Module
port 14: -- N2 No_Module
```

Migrating Storage on HP-UX Servers To MDS Switches

```
port 15: -- N2 No_Module
Brocade:admin>
```

The switchShow command reports the Domain_ID of the switch. In this case it can be seen that the Brocade switch has the Domain_ID set to 4.

```
Brocade:admin> nsShow
{
  Type Pid      COS  PortName                               NodeName                               TTL (sec)
  N   040000; 2,3; 50:06:04:82:bf:d1:db:cd; 50:06:04:82:bf:d1:db:cd; na FC4s:
FCP [EMC SYMMETRIX 5568] Fabric Port Name: 20:00:00:60:69:51:3a:60
  N   040100;      3;50:06:0b:00:00:11:b9:18;50:06:0b:00:00:11:b9:19; na FC4s:
FCP Fabric Port Name: 20:01:00:60:69:51:3a:60

The Local Name Server has 2 entries}
```

The **nsShow** reports the FC_ID of the port that have logged into to the switch. The Brocade command line interface reports this as port id or PID (in short). From the above output the FC_ID / PID of the storage port is 0x040000 (in hex).

At this juncture it is safe to remove the cables connected to the Brocade from both the storage (target) and the host as shown in Figure 20.

4.2.2 Connect path through MDS1

Before connecting the cables through the switch MDS1, the following configurations need to be done. On MDS1 set up a separate VSAN and set the Domain_ID to 4. Then configure the FC_ID of the storage port to the same value as it was on the Brocade switch. In our example, the FC_ID of the storage port (target) on the Brocade switch was 0x040000. When the storage (target) port is connected through MDS1 the statically assigned FC_ID is given to the storage (target) port. This step is very critical if the same CTD address for the disks have to be maintained on the host after the switch migration. This is needed to maintain the same Hardware path as explained in section 1.




Figure 21: MDS1 in the fabric

Migrating Storage on HP-UX Servers To MDS Switches

4.2.3 MDS1 configuration steps

```
MDS1(config-vsan-db)# vsan 40 name hpMDS
MDS1(config-vsan-db)# vsan 40 interface fc 3/8 ← Adding ports to VSAN 40
MDS1(config-vsan-db)# vsan 40 interface fc 8/25 ← Adding ports to VSAN 40
MDS1(config-vsan-db)# end
MDS1# show fcdomain vsan 40
The local switch is the Principal Switch.
```

Local switch run time information:

```
State: Stable
Local switch WWN: 20:28:00:05:30:00:63:5f
Running fabric name: 20:28:00:05:30:00:63:5f
Running priority: 128
Current Domain_ID: 0x62(98) ← Switch assigned Domain_ID
```

Local switch configuration information:

```
State: Enabled
FCID persistence: Disabled
Auto-reconfiguration: Disabled
Contiguous-allocation: Disabled
Configured fabric name: 20:01:00:05:30:00:28:df
Configured priority: 128
Configured Domain_ID: 0x00(0) (preferred)
```

Principal switch run time information:

```
Running priority: 128
```

No interfaces available.

The switch assigned Domain_ID for VSAN 40 is 98. It must be set to 4. This process is only disruptive to the new VSAN 40. All other VSANs on the switch are not affected.

```
MDS1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
MDS1(config)# fcdomain domain 4 static vsan 40 ← This sets the Domain_ID to 4
MDS1(config)# fcdomain restart disruptive vsan 40 ←restarts VSAN w/ Domain_ID 4
MDS1(config)# fcdomain fcid persistent vsan 40
MDS1(config)# fcdomain fcid database
MDS1(config-fcid-db)# vsan 40 wwn 50:06:04:82:bf:d1:db:cd fcid 0x040000 dynamic
                                                ← assign FC_ID to target
MDS1(config-fcid-db)#end
MDS1#
```

Now connect the target (storage) port and the host HBA to the switch ports that were made a part of VSAN 40. Zone the target and the HBA so that the HBA can see storage. The storage port gets the statically assigned FC_ID 0x040000. The host HBA's FC_ID is dynamically assigned. This ensures that the Hardware path for the disks does not change.

Issuing the **ioscan -fnC disk** command on the host rescans the bus and adds back the same devices. Now all should be the same as before except that one of the paths is through MDS1 as show in Figure 21. The above process is non disruptive.

4.3 Migrating the ports on McData to MDS

In this phase the ports connected to the McData need to be migrated over to the switch MDS2.

Migrating Storage on HP-UX Servers To MDS Switches

4.3.1 Gather data before removing the path through McData switch

Before getting started with this phase of migration, information regarding the various disks needs to be gathered. The **ioscan** and **vgdisplay** commands are used to gather the required information.

```
sjhp2:/# ioscan -fnkC disk
Class      I  H/W Path          Driver    S/W State  H/W Type  Description
=====
disk       0  0/0/2/0.0.0        sdisk     CLAIMED    DEVICE     SEAGATE
ST318404LC /dev/dsk/c2t0d0    /dev/rdisk/c2t0d0
disk       1  0/0/2/1.2.0        sdisk     CLAIMED    DEVICE     HP         DVD-ROM
304        /dev/dsk/c3t2d0    /dev/rdisk/c3t2d0
disk       2  0/1/1/0.97.4.19.0.0.0 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c4t0d0    /dev/rdisk/c4t0d0
disk       3  0/1/1/0.97.4.19.0.0.1 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c4t0d1    /dev/rdisk/c4t0d1
disk       4  0/1/1/0.97.4.19.0.0.2 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c4t0d2    /dev/rdisk/c4t0d2
disk       5  0/1/1/0.97.4.19.0.0.3 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c4t0d3    /dev/rdisk/c4t0d3
disk       6  0/1/1/0.97.4.19.0.0.4 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c4t0d4    /dev/rdisk/c4t0d4
disk       7  0/1/1/0.97.4.19.0.0.5 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c4t0d5    /dev/rdisk/c4t0d5
disk       8  0/1/3/0.4.0.0.0.0.0 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d0    /dev/rdisk/c6t0d0
disk       9  0/1/3/0.4.0.0.0.0.1 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d1    /dev/rdisk/c6t0d1
disk      10  0/1/3/0.4.0.0.0.0.2 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d2    /dev/rdisk/c6t0d2
disk      11  0/1/3/0.4.0.0.0.0.3 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d3    /dev/rdisk/c6t0d3
disk      12  0/1/3/0.4.0.0.0.0.4 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d4    /dev/rdisk/c6t0d4
disk      13  0/1/3/0.4.0.0.0.0.5 sdisk     CLAIMED    DEVICE     EMC
SYMMETRIX  /dev/dsk/c6t0d5    /dev/rdisk/c6t0d5
```

Based on the ioscan output, based on the FC_ID portion (highlighted in blue) of the Hardware path the Hardware path for disks c4t0d0 ... d7 is the path through the McData switch (Domain_ID 1). Similarly the Hardware path for disks c6t0d0 ... d7 is the path through the MDS1 switch (Domain_ID 4).

```
sjhp2:/dev/vghpMDS# vgdisplay -v vghpMDS
--- Volume groups ---
VG Name                /dev/vghpMDS
VG Write Access        read/write
VG Status               available
Max LV                 255
Cur LV                 3
Open LV                3
Max PV                 16
Cur PV                 5
Act PV                 5
Max PE per PV          2157
VGDA                   10
PE Size (Mbytes)       4
Total PE               10785
Alloc PE                7500
Free PE                3285
```

Migrating Storage on HP-UX Servers To MDS Switches

Total PVG	0
Total Spare PVs	0
Total Spare PVs in use	0

--- Logical volumes ---

LV Name	/dev/vghpMDS/lvol1
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/lvol2
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/lvol3
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1
PV Name	/dev/dsk/c6t0d1 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d2
PV Name	/dev/dsk/c6t0d2 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d3
PV Name	/dev/dsk/c6t0d3 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d4
PV Name	/dev/dsk/c6t0d4 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

PV Name	/dev/dsk/c4t0d5
PV Name	/dev/dsk/c6t0d5 Alternate Link
PV Status	available
Total PE	2157
Free PE	657
Autoswitch	On

Migrating Storage on HP-UX Servers To MDS Switches

The above vgdisplay output show that the disks c4t0d0 ... d7 belong to the primary path which is active and c6t0d0 ... d7 belong to the alternate path which is passive. In the next step the disks c6t0d1-d5 are removed from the disk group vghpMDS in preparation for the migration to MDS switch.

Before removing the primary path it is recommended to fail the path over to the alternate path to avoid any disruption of I/O to the host. This is achieved by disabling the active path using the vgreduce command. This caused the system to fail over to the alternate path.

The path through the McData needs to be migrated over to MDS2. Before migration of the path the FC_ID and the Domain_ID information from the McData are needed. This can be gathered on the McData switch with the **show switch** and **show name server** commands at the Root> prompt.

```
Root>show switch
Switch Information
State:                               Online
BB Credit:                           16
R_A_TOV:                             100
E_D_TOV:                             20
Preferred Domain_ID:                 1          ← Preferred Domain_ID
Switch Priority:                      Default
Speed:                               2 Gb/sec
Rerouting Delay:                     Enabled
Operating Mode:                      Open Systems
Interop Mode:                        McDATA Fabric 1.0
Active Domain_ID:                    1          ← Actual Domain_ID
World Wide Name:                     10:00:08:00:88:A0:D2:73
```

The **show switch** command reports the Domain_ID of the switch. In this case the McData switch has the Domain_ID 1.

```
Root> show nameServer
Type  Port Id Port Name                               Node Name                               COS  FC4 Types
-----
fPort 610413  50:06:04:82:BF:D1:DB:D2  50:06:04:82:BF:D1:DB:D2  2-3   2,18
fPort 610513  50:06:0B:00:00:11:BA:9B  50:06:0B:00:00:11:BA:9A   3     2,12
```

The “**show nameServer**” reports the FC_ID of the ports that have logged into to the switch. From the above output the FC_ID of the storage port is 0x610013 (in hex). In the case of McData even though the Domain_ID is set to 1 it actually shows up as 61 (hex) or 97 (decimal). The reason for this is explained in

4.3.2 Remove the path through McData switch

The active path through the McData needs to be removed. Once this is done the I/O fails over to the path through MDS1. **vgreduce** is used to remove this path from the volume group.

```
sjhp2:/# vgreduce /dev/vghpMDS /dev/dsk/c4t0d1 /dev/dsk/c4t0d2 /dev/dsk/c4t0d3
/dev/dsk/c4t0d4 /dev/dsk/c4t0d5
Device file path "/dev/dsk/c4t0d1" is an primary link.
Removing primary link and switching to an alternate link.
Device file path "/dev/dsk/c4t0d2" is an primary link.
Removing primary link and switching to an alternate link.
Device file path "/dev/dsk/c4t0d3" is an primary link.
Removing primary link and switching to an alternate link.
Device file path "/dev/dsk/c4t0d4" is an primary link.
Removing primary link and switching to an alternate link.
Device file path "/dev/dsk/c4t0d5" is an primary link.
Removing primary link and switching to an alternate link.
Volume group "/dev/vghpMDS" has been successfully reduced.
Volume Group configuration for /dev/vghpMDS has been saved in
/etc/lvmconf/vghpMDS.conf
```




Figure 22 : Disconnect McData Switch form topology

At this juncture it is safe to remove the cables connected to the McData switch from both the storage (target) and the host as shown in Figure 22.

4.3.3 Connect path through MDS2

Before connecting the cables through the MDS2 the following configurations need to be done. On MDS2 set up a separate VSAN and set the Domain_ID to 97. Then configure the FC_ID of the storage port to the same value as it had on the McData switch. The FC_ID of the storage port (target) on the McData was 0x610013.

When the storage (target) port is connected through MDS2 the statically assigned FC_ID is given to the storage (target) port. This step is very critical if the same CTD address for

Migrating Storage on HP-UX Servers To MDS Switches

the disks have to be maintained on the host after the switch migration. As a result, the Hardware path remains unchanged as explained in section 1.

The final topology after the migration to the two MDS switches is shown in Figure 23.




Figure 23: All MDS Topology (Migration 2)

```
MDS2(config-vsan-db)# vsan 50 name hpMDS
MDS2(config-vsan-db)# vsan 50 interface fc 3/8 ← Adding ports to VSAN 50
MDS2(config-vsan-db)# vsan 50 interface fc 8/25 ← Adding ports to VSAN 50
MDS2(config-vsan-db)# end
```

```
MDS2# show fcdomain vsan 50
```

The local switch is the Principal Switch.

Local switch run time information:

```
State: Stable
Local switch WWN: 20:32:00:05:30:00:63:9f
Running fabric name: 20:32:00:05:30:00:63:9f
Running priority: 128
Current Domain_ID: 0x76(118) ← Switch assigned Domain_ID
```

Local switch configuration information:

```
State: Enabled
FCID persistence: Disabled
Auto-reconfiguration: Disabled
Contiguous-allocation: Disabled
Configured fabric name: 20:01:00:05:30:00:28:df
Configured priority: 128
Configured Domain_ID: 0x00(0) (preferred)
```

Principal switch run time information:

```
Running priority: 128
```

No interfaces available.

The switch assigned Domain_ID for VSAN 50 is 118. It must be set to 97. This process is only disruptive to the new VSAN 50. All other VSANs on the switch are not affected.

Migrating Storage on HP-UX Servers To MDS Switches

```
MDS2# conf t
Enter configuration commands, one per line. End with CNTL/Z.
MDS2(config)# fcdomain domain 97 static vsan 50
MDS2(config)# fcdomain restart disruptive vsan 50
MDS2(config)# fcdomain fcid persistent vsan 50
MDS2(config)# fcdomain fcid database
MDS2(config-fcid-db)# vsan 50 wwn 50:06:04:82:bf:d1:db:d2 fcid 0x610413 dynamic
MDS2(config-fcid-db)#end
MDS2#
```

Sets the Domain_ID to 97

Restarts VSAN with Domain_ID 97

Assigns FC_ID to target

Now connect the target (storage) port and the host HBA to the switch ports that were made a part of VSAN 50. Zone the target and the HBA so that the HBA can see storage. The storage port gets the statically assigned FC_ID 0x610413. The host HBA's FC_ID is also statically assigned. This ensures that the Hardware path for the disks does not change.

The command **ioscan -fnc disk** on the host rescans the bus and adds back the same devices. The old devices are added back into volume group using the **vgextend** command.

```
sjhp2:/# vgextend /dev/vghpMDS /dev/dsk/c4t0d1 /dev/dsk/c4t0d2 /dev/dsk/c4t0d3
/dev/dsk/c4t0d4 /dev/dsk/c4t0d5
Volume group "/dev/vghpMDS" has been successfully extended.
Volume Group configuration for /dev/vghpMDS has been saved in
/etc/lvmconf/vghpMDS.conf
```

vgdisplay -v vghpMDS shows all the disks in the volume group vghpMDs having alternate links.

At this time the path through the McData switch has been migrated over to the switch MDS2 as shown in Figure 23. The above process is non disruptive and the CTD address of the disks are same as they were before the migration to the MDS switches.

4.4 Pros and Cons

4.4.1 Pros

1. The whole migration is non disruptive.
2. The whole process can be done online (while the host and applications are still up and running)
3. Only one path needs to be migrated at a time. The alternate path is always available to carry the I/O when one of the paths is being migrated.
4. Fail back to the old infrastructure is relatively simple.
5. The **CTD** addresses of disks do not change.
6. ***This is the only method that can be used to migrate a 3rd party switch to MDS while the application is online on a MC Service Guard Cluster running Oracle parallel server non disruptively. This is possible because of the persistent FC_ID feature available on MDS switches.***

4.4.2 Cons

1. The migration is more involved. The FC_ID of the target (storage) port i.e. the Domain_ID, area and the port # on the old infrastructure must be duplicated on the new infrastructure to maintain the same CTD address for the disks.

Migrating Storage on HP-UX Servers To MDS Switches

disk	5	0/1/1/0.4.0.0.0.0.3	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d3	/dev/rdisk/c4t0d3			
disk	6	0/1/1/0.4.0.0.0.0.4	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d4	/dev/rdisk/c4t0d4			
disk	7	0/1/1/0.4.0.0.0.0.5	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d5	/dev/rdisk/c4t0d5			
disk	8	0/1/3/0.5.0.0.0.0.0	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d0	/dev/rdisk/c6t0d0			
disk	9	0/1/3/0.5.0.0.0.0.1	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d1	/dev/rdisk/c6t0d1			
disk	10	0/1/3/0.5.0.0.0.0.2	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d2	/dev/rdisk/c6t0d2			
disk	11	0/1/3/0.5.0.0.0.0.3	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d3	/dev/rdisk/c6t0d3			
disk	12	0/1/3/0.5.0.0.0.0.4	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d4	/dev/rdisk/c6t0d4			
disk	13	0/1/3/0.5.0.0.0.0.5	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c6t0d5	/dev/rdisk/c6t0d5			

The device path name indicates that the HBA on Hardware path 1/1/0 is connected to a Brocade switch with Domain_ID set to 4. Similarly the HBA on Hardware path 1/3/0 is connected to a Brocade with Domain_ID set to 5.

The command **vgdisplay -v** lists details about all volume groups on a particular host.

```
sjhp2:/# vgdisplay -v
--- Volume groups ---
VG Name                /dev/vg00
VG Write Access         read/write
VG Status               available
Max LV                 255
Cur LV                 8
Open LV                 8
Max PV                 16
Cur PV                 1
Act PV                  1
Max PE per PV          4350
VGDA                    2
PE Size (Mbytes)       4
Total PE                4340
Alloc PE                3056
Free PE                 1284
Total PVG                0
Total Spare PVs         0
Total Spare PVs in use  0

--- Logical volumes ---
LV Name                 /dev/vg00/lvol1
LV Status                available/syncd
LV Size (Mbytes)         300
Current LE                75
Allocated PE              75
Used PV                   1

LV Name                 /dev/vg00/lvol2
LV Status                available/syncd
LV Size (Mbytes)         4096
Current LE                1024
Allocated PE              1024
```

Migrating Storage on HP-UX Servers To MDS Switches

Used PV	1
LV Name	/dev/vg00/lvol3
LV Status	available/syncd
LV Size (Mbytes)	200
Current LE	50
Allocated PE	50
Used PV	1
LV Name	/dev/vg00/lvol4
LV Status	available/syncd
LV Size (Mbytes)	200
Current LE	50
Allocated PE	50
Used PV	1
LV Name	/dev/vg00/lvol5
LV Status	available/syncd
LV Size (Mbytes)	20
Current LE	5
Allocated PE	5
Used PV	1
LV Name	/dev/vg00/lvol6
LV Status	available/syncd
LV Size (Mbytes)	1112
Current LE	278
Allocated PE	278
Used PV	1
LV Name	/dev/vg00/lvol7
LV Status	available/syncd
LV Size (Mbytes)	1700
Current LE	425
Allocated PE	425
Used PV	1
LV Name	/dev/vg00/lvol8
LV Status	available/syncd
LV Size (Mbytes)	4596
Current LE	1149
Allocated PE	1149
Used PV	1
--- Physical volumes ---	
PV Name	/dev/dsk/c2t0d0
PV Status	available
Total PE	4340
Free PE	1284
Autoswitch	On
--- Volume groups ---	
VG Name	/dev/vghpMDS
VG Write Access	read/write
VG Status	available
Max LV	255
Cur LV	3
Open LV	3
Max PV	16
Cur PV	5
Act PV	5

Migrating Storage on HP-UX Servers To MDS Switches

Max PE per PV	2157
VGDA	10
PE Size (Mbytes)	4
Total PE	10785
Alloc PE	7500
Free PE	3285
Total PVG	0
Total Spare PVs	0
Total Spare PVs in use	0

--- Logical volumes ---

LV Name	/dev/vghpMDS/vol1
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol2
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

LV Name	/dev/vghpMDS/vol3
LV Status	available/syncd
LV Size (Mbytes)	10000
Current LE	2500
Allocated PE	2500
Used PV	5

--- Physical volumes ---

PV Name	/dev/dsk/c4t0d1	
PV Name	/dev/dsk/c6t0d1	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d2	
PV Name	/dev/dsk/c6t0d2	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d3	
PV Name	/dev/dsk/c6t0d3	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d4	
PV Name	/dev/dsk/c6t0d4	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	

PV Name	/dev/dsk/c4t0d5	
---------	-----------------	--

Migrating Storage on HP-UX Servers To MDS Switches

PV Name	/dev/dsk/c6t0d5	Alternate Link
PV Status	available	
Total PE	2157	
Free PE	657	
Autoswitch	On	
sjhp2:/#		

There are only two volume groups (vg00 and vghpMDS) on this host. vg00 is the root volume group which is on local disks. This volume group has only one disk c2t0d0 as seen from the **ioscan -fnkc disk** output.

```
disk 0 0/0/2/0.0.0 sdisk CLAIMED DEVICE SEAGATE ST318404LC /dev/dsk/c2t0d0
/dev/rdisk/c2t0d0
```

On this host there is only one volume group vghpMDS that is affected when the Brocade switches are replaced with MDS switches. The disks affected by the migration are c4t0d1-d5 and c6t0d0 ...d5.

5.1.2 Create export map files for the affected volume groups

The **vgexport** command is used to create the required map files. To create the map files the **vgexport** command is used with the preview option.

```
sjhp2:/# vgexport -m /tmp/vghpmds.map -p -v -f /tmp/physical_devs.out vghpMDS
Beginning the export process on Volume Group "vghpMDS".
vgexport: Volume group "vghpMDS" is still active.
/dev/dsk/c4t0d1
/dev/dsk/c4t0d2
/dev/dsk/c4t0d3
/dev/dsk/c4t0d4
/dev/dsk/c4t0d5
/dev/dsk/c6t0d1
/dev/dsk/c6t0d2
/dev/dsk/c6t0d3
/dev/dsk/c6t0d4
/dev/dsk/c6t0d5
sjhp2:/#
```

This command results in two files being generated. The file vghpMDS.map file contains the volume group and associated logical volumes information. The file physical_devs.out contains the physical devices in the volume group. These files must be preserved for latter use. The contents of these files are show below:

```
sjhp2:/tmp# cat /tmp/vghpmds.map
1 lvol1
2 lvol2
3 lvol3
sjhp2:#/

sjhp2:/tmp# cat /tmp/physical_devs.out
/dev/dsk/c4t0d1
/dev/dsk/c4t0d2
/dev/dsk/c4t0d3
/dev/dsk/c4t0d4
/dev/dsk/c4t0d5
/dev/dsk/c6t0d1
```

Migrating Storage on HP-UX Servers To MDS Switches

```
/dev/dsk/c6t0d2  
/dev/dsk/c6t0d3  
/dev/dsk/c6t0d4  
/dev/dsk/c6t0d5  
sjhp2:/tmp#
```

The man page explanation for the `vgexport` command reveals the following options:

- m option creates a file which contains a description of the volume group and its associated logical volume(s) (if any).
- p option runs the `vgexport` in preview mode and does not export the volume group
- v is verbose
- f option Write the current set of `pv_paths` for the volume group to outfile. The outfile may then be used as the infile for the `vgimport –f` command.

The application has to be stopped, and the file systems using LUNs on the target port have to be un-mounted. Once the filesystems have been un-mounted and the applications shutdown, the volume group can be deactivated prior to exporting the volume group. This is done using the **vgchange** command. The **vgchange -a n /dev/vghpMDS** makes the volume group `vghpMDS` unavailable on the host.

The volume group device file information in the `/dev/vghpMDS` directory needs to be saved as well. The minor number of the group device file is needed to import the volume group back.

Note: Major and minor numbers are associated with the device special files in the /dev directory. These are used by the operating system to determine the actual driver and device to be accessed by the user-level requests to access the devices.

```
sjhp2:/tmp# ls -sl /dev/vghpMDS/group  
0 crw-r--r-- 1 root sys 64 0x010000 Oct 2 17:52 /dev/vghpMDS/group
```

```
sjhp2:/tmp# vgchange -a n /dev/vghpMDS  
Volume group "/dev/vghpMDS" has been successfully changed.  
sjhp2:/tmp#
```

The export removes the volume group from the host completely.

```
sjhp2:/tmp# vgexport -v /dev/vghpMDS  
Beginning the export process on Volume Group "/dev/vghpMDS".  
/dev/dsk/c4t0d1  
/dev/dsk/c4t0d2  
/dev/dsk/c4t0d3  
/dev/dsk/c4t0d4  
/dev/dsk/c4t0d5  
/dev/dsk/c6t0d1  
/dev/dsk/c6t0d2  
/dev/dsk/c6t0d3  
/dev/dsk/c6t0d4  
/dev/dsk/c6t0d5  
Volume group "/dev/vghpMDS" has been successfully removed.  
sjhp2:/tmp#
```

Migrating Storage on HP-UX Servers To MDS Switches

This removes the volume group from the system. The cables can now be disconnected from the Brocade switches as shown in Figure 25.




Figure 25 : Disconnect the cables from the Brocade switches

Reconnect the cables through the MDS switches as shown in Figure 26.

5.1.3 Rescan the devices and import the volume group




Figure 26: All MDS topology after migration

Once the zoning is complete the **ioscan** and **insf** commands need to be run on the host to create new devices. First a directory with the same name (vghpMDS) needs to be created in the /dev/ directory on the host. Then the group device file needs to be created with the same major and minor number. The **vgimport** is then used to import the volume group. The **vgchange** command is used to make the disk group available on the host and then the various file systems can be mounted again on the host. The **vgimport** command requires two input files, a logical volume map file and a physical devices map. The logical volume map file is the vghpmds.map file, and the physical map file that

Migrating Storage on HP-UX Servers To MDS Switches

contains a list of the new disks that are scanned and created through **ioscan** and **insf** commands.

```
sjhp2:/tmp# ioscan -fnkC disk
```

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
disk	0	0/0/2/0.0.0	sdisk	CLAIMED	DEVICE	SEAGATE
ST318404LC		/dev/dsk/c2t0d0	/dev/rdisk/c2t0d0			
disk	1	0/0/2/1.2.0	sdisk	CLAIMED	DEVICE	HP DVD-ROM
304		/dev/dsk/c3t2d0	/dev/rdisk/c3t2d0			
disk	14	0/1/1/0.20.0.1.0.0.0	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c8t0d0	/dev/rdisk/c8t0d0			
disk	15	0/1/1/0.20.0.1.0.0.1	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c8t0d1	/dev/rdisk/c8t0d1			
disk	16	0/1/1/0.20.0.1.0.0.2	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c8t0d2	/dev/rdisk/c8t0d2			
disk	17	0/1/1/0.20.0.1.0.0.3	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c8t0d3	/dev/rdisk/c8t0d3			
disk	18	0/1/1/0.20.0.1.0.0.4	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c8t0d4	/dev/rdisk/c8t0d4			
disk	19	0/1/1/0.20.0.1.0.0.5	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c8t0d5	/dev/rdisk/c8t0d5			
disk	20	0/1/3/0.66.0.1.0.0.0	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c9t0d0	/dev/rdisk/c9t0d0			
disk	21	0/1/3/0.66.0.1.0.0.1	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c9t0d1	/dev/rdisk/c9t0d1			
disk	22	0/1/3/0.66.0.1.0.0.2	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c9t0d2	/dev/rdisk/c9t0d2			
disk	23	0/1/3/0.66.0.1.0.0.3	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c9t0d3	/dev/rdisk/c9t0d3			
disk	24	0/1/3/0.66.0.1.0.0.4	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c9t0d4	/dev/rdisk/c9t0d4			
disk	25	0/1/3/0.66.0.1.0.0.5	sdisk	CLAIMED	DEVICE	EMC
SYMMETRIX		/dev/dsk/c9t0d5	/dev/rdisk/c9t0d5			

```
sjhp2:/tmp#
```

From the **ioscan -fnkC disk** output the newly created devices are c8t0d0-d5 and c9t0d1-d5. The Domain_IDs on the Cisco switches are 20 and 66. The old devices do not show up if the server is rebooted. Otherwise the **rmsf** command is need to remove the defunct devices.

The vghpMDS directory needs to be created in the /dev tree and the character device file group has to created as shown below.

```
sjhp2:/tmp# mkdir /dev/vghpMDS
sjhp2:/tmp# mknod /dev/vghpMDS/group c 64 0x010000
```

The physical_devs.out file needs to be recreated using the new devices. In this above migration it is:

```
sjhp2:/tmp# cat /tmp/physical_devs.out
/dev/dsk/c8t0d1
/dev/dsk/c8t0d2
/dev/dsk/c8t0d3
/dev/dsk/c8t0d4
/dev/dsk/c8t0d5
/dev/dsk/c9t0d1
```

Migrating Storage on HP-UX Servers To MDS Switches

```
/dev/dsk/c9t0d2  
/dev/dsk/c9t0d3  
/dev/dsk/c9t0d4  
/dev/dsk/c9t0d5  
sjhp2:/tmp#
```

The **vgimport** command is used to import the volume group back on the host.

```
sjhp2:/tmp# vgimport -m /tmp/vghpmds.map -v -f /tmp/physical_devs.out vghpMDS  
Beginning the import process on Volume Group "vghpMDS".  
vgimport: Warning: Volume Group contains "5" PVs, "10" specified. Continuing.  
Logical volume "/dev/vghpMDS/lvol1" has been successfully created with lv  
number 1.  
Logical volume "/dev/vghpMDS/lvol2" has been successfully created with lv  
number 2.  
Logical volume "/dev/vghpMDS/lvol3" has been successfully created with lv  
number 3.  
Volume group "/dev/vghpMDS" has been successfully created.  
Warning: A backup of this volume group may not exist on this machine.  
Please remember to take a backup using the vgcfgbackup command after activating  
the volume group.  
sjhp2:/tmp#
```

The **vgchange** command issued to make the volume group available on the host and the various file systems are then mounted and the applications can be restarted.

```
sjhp2:/# # vgchange -a n /dev/vghpMDS  
Volume group "/dev/vghpMDS" has been successfully changed.  
sjhp2:/tmp#
```

The file systems can now be mounted and the application restarted. The above process is disruptive and the CTD address of the disks will be different from what they were before the migration to the MDS switch.

5.2 Pros and cons

5.2.1 Pros

- Consolidating many islands is easier as they can be collapsed into one VSAN.
- This method is used when the administrator wants to maintain same CTD (new) addresses on both sides of the cluster.

5.2.2 Cons

- ***This migration requires planned down time.***
- The volume groups have to be exported and later imported with new disks.
- For the import and export to be done, a map of the volume group must be created before they are exported.
- More up front work is needed such as identifying disk associations with the volume group and also the various alternate paths involved.
- A Volume group can also be imported without the map but this is not recommended method. The note below explains the vgscan process and its limitations.

*Note: The import of the volume group with the new disks can be done with out the map files. This is done using the **vgscan** utility in LVM. This process has some issues. The **vgscan** goes out and scans all physical disks attached to the server looking for Logical volumes. It groups these physical volumes into volume groups by matching the volume group information found on the physical volumes. Then it searches the /dev directory for all group device files with the LVM major number, and tries to match device files with the logical volumes' information found on the physical volumes. If matches occur, it determines the volume group name from the device file path, and updates the /etc/lvmtab file with the volume group name and the list of physical volumes paths contained in that volume group. For volume groups where the device files cannot be matched, it prints the list of physical volumes for each volume group. After **vgscan** completes successfully, **vgimport** is run against each set of unmatched physical volumes.*

*This process is not with out problems. On systems with a large number of disks (physical volumes) the vgscan command tends to time out before completing the scan of all the disks attached to the server. This is because it runs only for a limited amount of time per invocation. This is by design, but unfortunately makes it impossible to scan all the disks and get the logical volume information needed to import the volume group. The **vgscan** process is impacted by the I/O happening on the system.*

6 Appendix A

6.1 FC_ID, Brocade and McData switches and HP-UX Hardware path

Both Brocade and McData switches encode the FC_IDs assigned to Nx ports that connect to them in a different manner.

6.1.1 Examples: Hardware path through Brocade switches.

In the case of Brocade the port field of the FC_ID is always 0 when attached in fabric mode. An example of this is shown in **Error! Reference source not found..** The area field is used to identify the port on the switch to which the target device or HBA is connected.




Figure 27 : FC_ID from the Brocade example

In the older (1 GB) switches only 4 of the 8 bits are used to identify a port on the switch. The upper 4 bits was always set to 0001. In this format of addressing only 16 ports can be supported. As the upper the 4 bits are fixed to 0001 the lowest port number was 16, i.e. 0001 0000 (bin) = 16 (decimal). So in the switches using this format to address ports, the port numbers start from 16 and go up to 32, for a total of 16 ports.

To address large number ports on a switch Brocade uses the “core PID format” for addressing the ports on the switches. In this mode of addressing they use all the 8 bits to address the ports. Thus port 0 in this format would show up as 0.

The below output is taken from an “**ioscan -fnkc disk**” command from a HP-UX host connected to a storage array through a Brocade switch.

```
disk 5 0/0/10/0/0.2.23.0.0.3.7 sdisk CLAIMED DEVICE EMC SYMMETRIX
/dev/dsk/c4t3d7 /dev/rdisk/c4t3d7
```

The Hardware path in the above example is from a Brocade switch that used traditional port addressing; an offset of 16 is added to the port. The Hardware path (0/0/10/0) is the location of the HBA on the host, 2.23.0 is the FC_ID of the target port, and 0.3.7 is the bus, target and LUN address of the device (disk). The FC_ID portion of the Hardware path of the target port is shown in Figure 27.

Migrating Storage on HP-UX Servers To MDS Switches

In the above example the Domain_ID is 2, the area field is used for the port # and is equal to 23, and the port field is 0 on all Brocade switches indicating that it is fabric attached. If the target were loop attached then the port field would be the AL_PA assigned to the port.

Since this switch uses the traditional port addressing the actual port is $23 - 16 = 7$.

```
disk 2 0/1/1/0.4.0.0.0.0.0 sdisk CLAIMED DEVICE EMC SYMMETRIX
/dev/dsk/c6t0d0 /dev/rdisk/c4t0d0
```

In the above example 0/1/1/0 is the HBA location on the host, 4.0.0 is the FC_ID of the target port and 0.0.0 is the bus, target and LUN address of the device (disk). This switch uses the core PID format for addressing the port. In FC_ID example 4.0.0, 4 is the Domain_ID of the switch 0 is the port # and 0 on Brocade indicates that it is fabric attached.

6.1.2 Examples: Hardware path through McData switches.

In the case of McData the first 3 bits are reserved and always set to **011** (hex) or 6 (decimal). So a maximum of $2^5 = 32$ unique Domain_IDs are possible with McData switches. McData supports Domain_IDs 1 – 31. Thus first Domain_ID would be **0110 0001**, this is 61(hex) or 97 (decimal). The Domain_IDs in hex on McData ranges from 0x61 to 0x7F. Thus the Domain_ID 1 shows up as 97 on the host Hardware field. McData uses 19 in the port field of the FC_ID to indicate that the storage (target) is fabric attached. McData port numbering starts from 4. An example of a McData switch is shown in Figure 28.

The below output is an “**ioscan -fnkc disk**” from a HP-UX host connected to a storage array through a McData switch.

```
disk 150 8/8/1/0.118.23.19.0.14.2 sdisk CLAIMED DEVICE EMC SYMMETRIX
/dev/dsk/c14t14d2 /dev/rdisk/c14t14d2
```

In the above example 8/8/1/0 is the Hardware location of HBA on host 118.23.19 is the FC_ID of the target port and 0.14.2 is the bus, target and LUN address. Further looking into the FC_ID portion of the address 118.23.19, it can be seen that the Domain_ID is 118, area bits are used to identify port and the port bits are always set to 19, indicating a fabric connection on McData switches.

Migrating Storage on HP-UX Servers To MDS Switches




Figure 28 : FC_ID from the McData Example

In case of a McData switch the actual Domain_ID which is Domain_ID – 96 (offset) as explained above would be $118 - 96 = 22$. The port Id is $23 - 4$ (numbering start from 4) = 17. Also the port bits are always set to 19 indicating that the port is fabric attached. (only McData)

```
disk      3  0/1/1/0.97.4.19.0.0.1  sdisk  CLAIMED  DEVICE  EMC  SYMMETRIX
/dev/dsk/c4t0d1  /dev/rdisk/c4t0d1
```

In the above example 0/1/1/0 is the HBA location on the host, 97.4.19 is the FC_ID of the target port and 0.0.1 is the bus, target and LUN address of the device (disk). The actual Domain_ID is $97 - 96 = 1$. The port # is $4 - 4 = 0$ and the last field is 19 which indicate that this port is fabric attached. (this is specific to McData switches only)

7 Appendix B

7.1 Veritas Volume Manager

Veritas Volume manager is also now available on HP-UX. If Veritas Volume manager is used the previously discussed problems are not encountered. This is because Veritas Volume Manager can accommodate the disk name changes as it scans the private region on each disk that it controls to recreate the volume group even if the disk name changes.

The disk layout of a Veritas Volume Manager disk is show in Figure 29. It consists of private region and the public region.




Figure 29: VxVM disk layout

7.1.1 VxVM basics

Veritas Volume Manager initializes a disk and organizes the disk into two regions the private region and the public region. The private region is a partition in a disk where the Volume Manager stores the data related to disk and their relationship to the volume group name, volume group ID, disk names, sub disks and plexes configuration, etc. The public region is where the user data is stored. To create a usable volume the public region is carved into smaller sub disks (if disk striping is used). Then one or more sub disks are then assembled together to form a plex then this plex is enclosed as a volume. This volume is then presented to the OS for use. At least one plex has to be associated with a volume. More than one plex can be associated with a volume. This is usually used to mirror data on to another plex. A file system can be then created on these volumes and then mounted on the system for use. Alternatively the volumes can be used as a raw volume as would be case with databases. The Figure 30 illustrates this process.

At system start up or when an explicit volume related commands are issued, the volume manager software scans all the disks visible to the system for disk group related configuration information in their respective private regions. The Volume Manager then imports all the volume groups if all the required disks needed are visible to the hosts.

Migrating Storage on HP-UX Servers To MDS Switches

Note: The ability of the volume manager to scan the private region dynamically recreate the volume group without having to depend on a static file for disk to volume group mapping and its ability to recognize and add the alternate path automatically, make the problems associated with the LVM related switch migration a non issue when using Veritas Volume Manager.




Figure 30: VxVM volume creation

7.1.2 Multi path with Veritas Volume Manager

DMP (Dynamic Multi path) is the multi-pathing software that comes with Veritas Volume Manager. DMP is an active / active multi-pathing software, which means that it can actively send traffic through all the paths available to a disk sub system. DMP does this using a round robin method. DMP can automatically detect alternate paths to the same disk. No manual configuration is needed. Failover of a disabled path is automatic.