



# Configuring Solaris10 x86 for Xyratex Storage Using MPxIO with ZFS

By David Lawler



## Notices

The information in this document is subject to change without notice.

While every effort has been made to ensure that all information in this document is accurate, Xyratex accepts no liability for any errors that may arise.

© 2008 Xyratex (the trading name of Xyratex Technology Limited). Registered Office: Langstone Road, Havant, Hampshire, PO9 1SA, England. Registered number 03134912.

No part of this document may be transmitted or copied in any form, or by any means, for any purpose, without the written permission of Xyratex.

Xyratex is a trademark of Xyratex Technology Limited. All other brand and product names are registered marks of their respective proprietors.

Issue 1.0 | November, 2008

# Contents

Introduction .....	2
Configuration Steps.....	2
Configure LUNs so They Are Visible to Solaris .....	2
Enabling/Disabling Multipathing .....	3
Edit The/kernel/drv/scsi_vhci.conf File.....	4
Perform a Reconfiguration Boot.....	4
Add LUNs to a zpool.....	6
Confirm Filesystem.....	7
Confirm Paths Using luxadm and cfgadm .....	7
Confirm Paths Using mpathadm .....	11
Failover/Failback .....	17
SAN Mapping Information .....	17
Performance .....	18

## Introduction

This document describes how to implement Sun Storage Traffic Manager (STMS a.k.a. MPxIO) path manager to allocate Xyratex LUNs across 2 controllers. Once MPxIO is configured as described below, then both controllers of a Xyratex Storage Device will be utilized. As in all multipath implementations MPxIO can not be used for load balancing of I/O to a LUN across the 2 controllers. BOTH controllers are in an active state, and as a result, the MPxIO is load balancing multiple LUNS across controllers as opposed to a single LUN across both controllers.

### Caveats

Minimum level for firmware is 2.4 (2.4.1 for F5404E).

Minimum of Solaris 10 Update 3 with SAN packages and all required patches.

### x86 Configuration

This document was created using Solaris x86 10 Update 6 installed with ZFS and connected to a F5404E.

## Configuration Steps

### Configure LUNs so They Are Visible to Solaris

```
root@glendry # format
```

```
Searching for disks...done
```

```
c1t1d0: configured with capacity of 33.92GB
```

```
c2t21000050CCF00E4Ad1: configured with capacity of 9.31GB
```

```
c3t23000050CCF00E4Ad1: configured with capacity of 9.31GB
```

### AVAILABLE DISK SELECTIONS

0. c1t0d0 <DEFAULT cyl 4424 alt 2 hd 255 sec 63>/pci@4,0/pci1028,14a@5/sd@0,0
1. c1t1d0 <SEAGATE-ST336607LC-DS09 cyl 49395 alt 2 hd 2 sec 720>/pci@4,0/pci1028,14a@5/sd@1,0
2. c2t21000050CCF00E4Ad0 <XYRATEX-F5404E- -9.32GB>/pci@1,0/pci10df,fa00@4/fp@0,0/disk@w21000050ccf00e4a,0
3. c2t21000050CCF00E4Ad1 <XYRATEX-F5404E- cyl 2383 alt 2 hd 64 sec 128>/pci@1,0/pci10df,fa00@4/fp@0,0/disk@w21000050ccf00e4a,1
4. c2t21000050CCF00E4Ad2 <DEFAULT cyl 1214 alt 2 hd 255 sec 63>/pci@1,0/pci10df,fa00@4/fp@0,0/disk@w21000050ccf00e4a,2
5. c3t23000050CCF00E4Ad0 <XYRATEX-F5404E- -9.32GB>/pci@1,0/pci10df,fa00@4,1/fp@0,0/disk@w23000050ccf00e4a,0
6. c3t23000050CCF00E4Ad1 <XYRATEX-F5404E- cyl 2383 alt 2 hd 64 sec 128>/pci@1,0/pci10df,fa00@4,1/fp@0,0/disk@w23000050ccf00e4a,1

```
7. c3t23000050CCF00E4Ad2 <DEFAULT cyl 1214 alt 2 hd 255 sec 63>/pci@1,0/pci10df,fa00@4,1/
   fp@0,0/disk@w23000050ccf00e4a,2
```

Specify disk (enter its number):

**NOTE: 3 x 10GB LUNs seen twice.**

## Enabling/Disabling Multipathing

Multipathing is enabled by default on x86 based systems. Use the following steps to disable multipathing or to enable multipathing if you have disabled it previously.

Caution — If the critical file systems needed for booting, such as the root (the `/`), `/usr`, or `/var` file systems, are on fibre channel disks, you should not disable multipathing on them otherwise your system will not boot. You may disable multipathing on all fibre channel devices if the critical file systems are not on fibre channel disks. If the critical file systems are on fibre channel disks, you may disable multipathing only on those fibre channel HBA (Host Bus Adapter) controller ports which aren't connected to the disks containing the critical file systems using the instructions in the section Enabling or Disabling Multipathing on a per-port basis.

- A. Open the `/kernel/drv/fp.conf` file in a text editor.
- B. If you want to enable multipathing — Change the value of `mpxio-disable` to “no”. If the entry doesn't exist add a new entry.

```
mpxio-disable="no";
```

If you want to disable multipathing — Change the value of `mpxio-disable` to “yes”. If the entry doesn't exist add a new entry.

```
mpxio-disable="yes";
```

**Note — This will not affect your per-port multipathing configuration. For information on per-port configuration refer to the Sun Microsystems document 819-0139 section titled “Enabling or Disabling Multipathing on a Per-Port Basis.”**

- c. Save and exit the `fp.conf` file.
- D. Perform a reconfiguration reboot.

```
# touch /reconfigure
# shutdown -g0 -y -i6
```

## Edit The/kernel/drv/scsi\_vhci.conf File

Multipathing can be configured on Xyratex symmetric storage devices. A symmetric storage device is defined as a storage device in which all paths to the storage device are active and I/O commands can be issued through any path.

Set load-balance to none (default in Round Robin).

Set auto-failback to enable (default is enable).

Set device-type-scsi-options-list to appropriate entry based on attached storage.

### \*Device-type-scsi-options-list entry

The vendor ID (identification) and product ID are the vendor and product identification strings that the device returns in SCSI INQUIRY data. The vendor ID must be 8 characters long. You must specify all 8 characters (even if the trailing characters are spaces). The product ID can be up to 16 characters long.

The information needed to populate this portion can be obtained by using the inquiry option to the format command.

### Example entries

```
load-balance="none";
auto-failback="enable";
device-type-scsi-options-list =
"XYRATEX F5402E", "symmetric-option",
"XYRATEX F5412E", "symmetric-option",
"XYRATEX F5404E", "symmetric-option";
symmetric-option = 0x1000000;
```

For further explanation of the entries in this file, see the actual /kernel/drv/scsi\_vhci.conf file for descriptions of the options.

## Perform a Reconfiguration Boot

```
# touch /reconfigure
```

Then reboot. Once the reboot is complete, confirm the devices and paths are recognized as being STMS managed devices.

Format will now show a single device instance with the scsi\_vhci naming convention.

```
root@glendry # format
```

```
Searching for disks...done
```



c1t1d0: configured with capacity of 33.92GB

c7t60050CC000F00E4A00000000000022Ad0: configured with capacity of 9.31GB

#### AVAILABLE DISK SELECTIONS

0. c1t0d0 <DEFAULT cyl 4424 alt 2 hd 255 sec 63>/pci@4,0/pci1028,14a@5/sd@0,0
1. c1t1d0 <SEAGATE-ST336607LC-DS09 cyl 49395 alt 2 hd 2 sec 720>/pci@4,0/pci1028,14a@5/sd@1,0
2. c7t60050CC000F00E4A00000000000022Ad0 <XYRATEX-F5404E-cyl2383alt2hd64sec 128>/scsi\_vhci/disk@g60050cc000f00e4a00000000000022a
3. c7t60050CC000F00E4A00000000000022Bd0 <DEFAULT cyl 1214 alt 2 hd 255 sec 63>/scsi\_vhci/disk@g60050cc000f00e4a00000000000022b
4. c7t60050CC000F00E4A000000000000229d0 <XYRATEX-F5404E- -9.32GB>/scsi\_vhci/disk@g60050cc000f00e4a000000000000229

Specify disk (enter its number):

#### Note — 3 x 10GB LUNS seen once

In the /dev and /devices trees, multipath-enabled devices receive new names that indicate that they are under MPxIO control. A previously configured device therefore will have a different name from its original name when it is under multipath control.

Device name with multipath disabled:

/dev/dsk/c1t1d0s0

Device name with multipath enabled:

/dev/dsk/c0t60003BA27D5170003E5D2A7A0007F3D2d0s0

Because of this, if you have applications that utilize device names, they must be configured to use the new names whenever you change multipath configuration from disabled to enabled or vice-versa. The system's /etc/vfstab file and the dump configuration also contain references to device names. Manually update the /etc/vfstab and the dump configuration with the new device names. You need to do this only if the device names listed in /etc/vfstab or your dump device are affected by the device name changes. Use dumpadm(1M) command to update the dump configuration.

## Add LUNs to a zpool

```
root@glendry # zpool create sumo c7t60050CC000F00E4A00000000000022Ad0 c7t60050CC000F00E-4A00000000000022Bd0 c7t60050CC000F00E4A000000000000229d0
```

```
root@glendry # zpool list
```

NAME	SIZE	USED	AVAIL	CAP	HEALTH	ALTROOT
glenfs	33.8Gb	6.43Gb	27.3Gb	19%	Online	—
sumo	27.8Gb	155K	27.7Gb	0%	Online	—

```
root@glendry # format
```

```
Searching for disks...done
```

```
c1t1d0: configured with capacity of 33.92GB
```

### AVAILABLE DISK SELECTIONS

0. c1t0d0 <DEFAULT cyl 4424 alt 2 hd 255 sec 63>/pci@4,0/pci1028,14a@5/sd@0,0
1. c1t1d0 <SEAGATE-ST336607LC-DS09 cyl 49395 alt 2 hd 2 sec 720>/pci@4,0/pci1028,14a@5/sd@1,0
2. c7t60050CC000F00E4A00000000000022Ad0 <XYRATEX-F5404E- -9.32GB>/scsi\_vhci/disk@g60050cc000f00e4a00000000000022a
3. c7t60050CC000F00E4A00000000000022Bd0 <XYRATEX-F5404E- -9.32GB>/scsi\_vhci/disk@g60050cc000f00e4a00000000000022b
4. c7t60050CC000F00E4A000000000000229d0 <XYRATEX-F5404E- -9.32GB>/scsi\_vhci/disk@g60050cc000f00e4a000000000000229

```
Specify disk (enter its number):
```

**What to look for** — All LUNs have same vendor and product ID. You will also want to know which LUN is associated with which StorView™ LD (Logical Drive). To do this, look at the last digit of the target in the logical drive number. It is in HEX and the lowest number is LD0 (“9” and drive number 4 in this case.) LD1 is “A” and LD2 is “B”. Normally these will be assigned in order and drive 2 would be LD0, however, in order to illustrate how this works, these were created out of order.

## Confirm Filesystem

Filesystem	KB	Used	Avail	Capacity	Mounted on
glenfs/ROOT/s10x_u6wos_07b	34836480	8065803	25195163	25%	/
/devices	0	0	0	0%	/devices
ctfs	0	0	0	0%	/system/contract
proc	0	0	0	0%	/proc
mnttab	0	0	0	0%	/etc/mnttab
swap	896416	372	896044	1%	/etc/svc/volatile
objfs	0	0	0	0%	/system/object
sharefs	0	0	0	0%	/etc/dfs/sharetab
/usr/lib/libc/libc_hwcap1.so.1	33260967	8065803	25195163	25%	/lib/libc.so.1
fd	0	0	0	0%	/dev/fd
swap	896052	8	896044	1%	/tmp
swap	896076	32	896044	1%	/var/run
glenfs/export	34836480	19	25195163	1%	/export
glenfs/export/home	34836480	18	25195163	1%	/export/home
glenfs	3483680	35	25195163	1%	/glenfs
sumo	28643328	2349119	26294056	9%	/sumo

**What to look for** — Confirm the newly created ZFS filesystem was created and mounted, /sumo in this instance.

## Confirm Paths Using luxadm and cfgadm

The following command examples show methods of confirming the pathing configuration using the luxadm and cfgadm command.

```
root@glendry # luxadm probe
```

Found fibre channel device(s):

```
Node WWN:20000050ccf00e4a Device Type:Disk device
```

```
Logical Path:/dev/rdisk/c7t60050CC000F00E4A00000000000022Bd0s2
```

```
Node WWN:20000050ccf00e4a Device Type:Disk device
```

```
Logical Path:/dev/rdisk/c7t60050CC000F00E4A00000000000022Ad0s2
```

```
Node WWN:20000050ccf00e4a Device Type:Disk device
```

```
Logical Path:/dev/rdisk/c7t60050CC000F00E4A000000000000229d0s2
```

**What this shows** — This provides the Configuration WWN for the storage device as seen in StorView (controller 0 & 1 > under “Configuration”) and the logical path for the LUNs, which are used in the following commands.

```
root@glendry # luxadm display /dev/rdisk/c7t60050CC000F00E4A00000000000022Bd0s2
```

```
DEVICE PROPERTIES for disk: /dev/rdisk/c7t60050CC000F00E4A00000000000022Bd0s2
```

```
Vendor                Xyratex
Product ID            SA 48
Revision
Serial Number         00B00E4A
Unformatted capacity  9540.000 MB
Write Cache           Enabled
Read Cache            Enabled
  Minimum prefetch    0x0
  Maximum prefetch    0xffff
Device Type           Disk device
Path(s)
```

```
/devices/scsi_vhci/disk@g60050cc000f00e4a00000000000022b:c,raw
```

```
/devices/scsi_vhci/disk@g60050cc000f00e4a00000000000022b:c,raw
```

```
Controller            /dev/cfg/c2
Device Address         21000050ccf00e4a,2
Host controller port WWN 10000000c93840a0
Class                  primary
State                  ONLINE
Controller            /dev/cfg/c3
Device Address         23000050ccf00e4a,2
Host controller port WWN 10000000c93840a1
Class                  primary
State                  ONLINE
```

```
root@glendry # luxadm display /dev/rdisk/c7t60050CC000F00E4A00000000000022Ad0s2
```

```
DEVICE PROPERTIES for disk: /dev/rdisk/c7t60050CC000F00E4A00000000000022Ad0s2
```

```
Vendor                Xyratex
Product ID            SA 48
Revision
Serial Number         00700E4A
Unformatted capacity  9540.000 MB
Write Cache           Enabled
Read Cache            Enabled
  Minimum prefetch    0x0
  Maximum prefetch    0xffff
Device Type           Disk device
Path(s)
```

/dev/rdsk/c7t60050CC000F00E4A00000000000022Ad0s2

/devices/scsi\_vhci/disk@g60050cc000f00e4a00000000000

0022a:c,raw

Controller	/dev/cfg/c2
Device Address	21000050ccf00e4a,1
Host controller port WWN	10000000c93840a0
Class	primary
State	ONLINE
Controller	/dev/cfg/c3
Device Address	23000050ccf00e4a,1
Host controller port WWN	10000000c93840a1
Class	primary
State	ONLINE

root@glendry # luxadm display /dev/rdsk/c7t60050CC000F00E4A0000000000000229d0s2

DEVICE PROPERTIES for disk: /dev/rdsk/c7t60050CC000F00E4A0000000000000229d0s2

Vendor	Xyratex
Product ID	SA 48
Revision	
Serial Number	00300E4A
Unformatted capacity	9540.000 MB
Write Cache	Enabled
Read Cache	Enabled
Minimum prefetch	0x0
Maximum prefetch	0xffff
Device Type	Disk device
Path(s)	

/dev/rdsk/c7t60050CC000F00E4A0000000000000229d0s2

/devices/scsi\_vhci/disk@g60050cc000f00e4a0000000000000229:c,raw

Controller	/dev/cfg/c2
Device Address	21000050ccf00e4a,0
Host controller port WWN	10000000c93840a0
Class	primary
State	ONLINE
Controller	/dev/cfg/c3
Device Address	23000050ccf00e4a,0
Host controller port WWN	10000000c93840a1
Class	primary
State	ONLINE

**What this shows** — The “State” of both controllers is “ONLINE” and the “Class” on both is “Primary.” When both controllers are listed as primary, you know the setting in the /kernel/drv/scsi\_vhci.conf file has a symmetric storage device listed after “device-type-scsi-options-list =” and the line “symmetric-option = 0x1000000;” exists and it uncommented.

This command also shows the vhci driver name for the LUN, in addition to it’s size, the host HBA WWN, as well as other helpful information.

```
root@glendry # cfgadm -al
Ap_Id                Type      Receptacle  Occupant    Condition
c1                   scsi-bus  connected   configured  unknown
c1::dsk/c1t0d0       disk      connected   configured  unknown
c1::dsk/c1t1d0       disk      connected   configured  unknown
c2                   fc-private connected   configured  unknown
c2::21000050ccf00e4a disk      connected   configured  unknown
c3                   fc-private connected   configured  unknown
c3::23000050ccf00e4a disk      connected   configured  unknown
```

**What to look for** — Confirm that both/all relevant controllers are connected and configured.

```
root@glendry # cfgadm -o show_FCP_dev -al
Ap_Id                Type      Receptacle  Occupant    Condition
c2                   fc-private connected   configured  unknown
c2::21000050ccf00e4a,0 disk      connected   configured  unknown
c2::21000050ccf00e4a,1 disk      connected   configured  unknown
c2::21000050ccf00e4a,2 disk      connected   configured  unknown
c2::21000050ccf00e4a,3 processor connected   configured  unknown
c3                   fc-private connected   configured  unknown
c3::23000050ccf00e4a,0 disk      connected   configured  unknown
c3::23000050ccf00e4a,1 disk      connected   configured  unknown
c3::23000050ccf00e4a,2 disk      connected   configured  unknown
c3::23000050ccf00e4a,3                connected   configured  unknown
```

**What to look for** — Confirm all LUNs are visible under each controller.

```
root@glendry # fcinfo hba-port
HBA Port WWN: 10000000c93840a0
    OS Device Name: /dev/cfg/c2
    Manufacturer: Emulex
    Model: LP10000DC
    Firmware Version: 1.90a4 (T2D1.90A4)
    FCode/BIOS Version: none
    Serial Number: MS34585577
    Driver Name: emlxs
```



```
Driver Version: 2.31h (2008.06.16.10.13)
disk
Type: L-port
State: online
Supported Speeds: 1Gb 2Gb
Current Speed: 2Gb
Node WWN: 20000000c93840a0
HBA Port WWN: 10000000c93840a1
OS Device Name: /dev/cfg/c3
Manufacturer: Emulex
Model: LP1000DC
Firmware Version: 1.90a4 (T2D1.90A4)
FCode/BIOS Version: none
Serial Number: MS34585577
Driver Name: emlxs
Driver Version: 2.31h (2008.06.16.10.13)
Type: L-port
State: online
Supported Speeds: 1Gb 2Gb
Current Speed: 2Gb
Node WWN: 20000000c93840a1
```

**What to look for** — This command provides excellent information about the HBA(s).

### Confirm Paths Using mpathadm

To confirm system is set up to support storage:

```
root@glendry # mpathadm show mpath-support libmpscsi_vhci.so
Vendor                XYRATEX
Product              F5404E
Revision
Supported Load Balance Types
                        round-robin
```

To show HBA and port:

```
root@glendry # mpathadm list initiator-port
Initiator Port      iqn.1986-03.com.sun:01:e0000000786b.4911f214,4000002a00ff
Initiator Port      10000000c93840a1
Initiator Port      10000000c93840a0
```

```
root@glendry # mpathadm show initiator-port 10000000c93840a1
Initiator Port:      10000000c93840a1
Transport Type:      Fibre Channel
OS Device File:      /devices/pci@1,0/pci10df,fa00@4,1/fp@0,0
```

```
root@glendry # mpathadm show initiator-port 10000000c93840a0
Initiator Port:      10000000c93840a0
Transport Type:      Fibre Channel
OS Device File:      /devices/pci@1,0/pci10df,fa00@4/fp@0,0
```

To show LUN information and path state:

```
root@glendry # mpathadm list lu
/dev/rdisk/c7t60050CC000F00E4A000000000000022Bd0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
    Total Path Count: 2
    Operational Path Count: 2
/dev/rdisk/c7t60050CC000F00E4A0000000000000229d0s2
    Total Path Count: 2
    Operational Path Count: 2
```

```
root@glendry # mpathadm show lu /dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
```

```
Logical Unit:      /dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
mpath-support:      libmpscsi_vhci.so
Vendor:             XYRATEX
Product:            F5404E
Revision:
Name Type:          unknown type
Name:               60050cc000f00e4a000000000000022a
Asymmetric:        no
Current Load Balance: none
Logical Unit Group ID: NA
Auto Failback:      on
Auto Probing:       NA
```

Paths:

```
Initiator Port Name: 10000000c93840a1
Target Port Name:    23000050ccf00e4a
Override Path:       NA
Path State:          OK
Disabled:            no
```

```

Initiator Port Name: 10000000c93840a0
Target Port Name: 21000050ccf00e4a
Override Path: NA
Path State: OK
Disabled: no

Target Ports:

Name: 23000050ccf00e4a
Relative ID: 0

Name: 21000050ccf00e4a
Relative ID: 0

```

To disable and/or enable a path:

```

root@glendry # mpathadm disable path -i 10000000c93840a1 -t 23000050ccf00e4a -l
/dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2

```

```

root@glendry # mpathadm show lu
/dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
mpath-support: libmpscsi_vhci.so
Vendor: XYRATEX
Product: F5404E
Revision:
Name Type: unknown type
Name: 60050cc000f00e4a0000000000000022a
Asymmetric: no
Current Load Balance: none
Logical Unit Group ID: NA
Auto Failback: on
Auto Probing: NA

```

Paths:

```

Initiator Port Name: 10000000c93840a1
Target Port Name: 23000050ccf00e4a
Override Path: NA
Path State: OK
Disabled: yes

```

```

Initiator Port Name: 10000000c93840a0
Target Port Name: 21000050ccf00e4a
Override Path: NA
Path State: OK
Disabled: no

```

Target Ports:

Name: 23000050ccf00e4a  
Relative ID: 0

Name: 21000050ccf00e4a  
Relative ID: 0

```
root@glendry # mpathadm enable path -i 10000000c93840a1 -t 23000050ccf00e4a -l /dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
```

```
root@glendry # mpathadm show lu
```

```
/dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
```

```
mpath-support: libmpscsi_vhci.so  
Vendor: XYRATEX  
Product: F5404E  
Revision:  
Name Type: unknown type  
Name: 60050cc000f00e4a000000000000022a  
Asymmetric: no  
Current Load Balance: none  
Logical Unit Group ID: NA  
Auto Failback: on  
Auto Probing: NA
```

Paths:

Initiator Port Name: 10000000c93840a1  
Target Port Name: 21000050ccf00e4a  
Override Path: NA  
Path State: OK  
Disabled: no

Initiator Port Name: 10000000c93840a0  
Target Port Name: 21000050ccf00e4a  
Override Path: NA  
Path State: OK  
Disabled: no

Target Ports:

Name: 23000050ccf00e4a  
Relative ID: 0

Name: 21000050ccf00e4a  
Relative ID: 0



To confirm both controllers are being utilized, start I/O to all LUNS. Open StorView and start “Logical Drive Statistics.” From here you can observe I/O running through each controller, and even more specifically, each logical drive.

If, with the above method, you find that all data is going through a single controller, do the following:

1) Determine which LUN(s) you need to adjust and then determine the HBA port and Controller information for the path you want the data to go through — i.e. to move LD1 off of controller 1 (c3) and onto controller 0 (c2), do the following:

1) Disable path.

```
root@glendry # mpathadm disable path -i 10000000c93840a1 -t 23000050ccf00e4a -l /dev/rdisk/c7t60050C-
C000F00E4A000000000000022Ad0s2
```

2) Confirm path is disabled.

```
root@glendry # mpathadm show lu /dev/rdisk/c7t60050CC000F00E4A00000000000022Ad0s2
```

```
Logical Unit:      /dev/rdisk/c7t60050CC000F00E4A00000000000022Ad0s2
```

```
  mpath-support:      libmptscsi_vhci.so
  Vendor:              XYRATEX
  Product:             F5404E
  Revision:
  Name Type:           unknown type
  Name:                60050cc000f00e4a000000000000022a
  Asymmetric:         no
  Current Load Balance: none
  Logical Unit Group ID: NA
  Auto Failback:      on
  Auto Probing:       NA
```

Paths:

```
  Initiator Port Name: 10000000c93840a1
  Target Port Name:   23000050ccf00e4a
  Override Path:      NA
  Path State:         OK
  Disabled:           yes
```

```
  Initiator Port Name: 10000000c93840a0
  Target Port Name:   21000050ccf00e4a
  Override Path:      NA
  Path State:         OK
  Disabled:           no
```

Target Ports:

Name:	23000050ccf00e4a
Relative ID:	0
Name:	21000050ccf00e4a
Relative ID:	0

3) Run I/O through both controllers. Use dd or cp.

4) With I/O running, re-enable the path.

```
root@glendry # mpathadm enable path -i 10000000c93840a1 -t 23000050ccf00e4a -l /dev/rdisk/c7t60050C-  
C000F00E4A000000000000022Ad0s2
```

5) Confirm the path is enabled.

```
root@glendry # mpathadm show lu dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2
```

```
Logical Unit: /dev/rdisk/c7t60050CC000F00E4A000000000000022Ad0s2  
mpath-support: libmptscsi_vhci.so  
Vendor: XYRATEX  
Product: F5404E  
Revision:  
Name Type: unknown type  
Name: 60050cc000f00e4a000000000000022a  
Asymmetric: no  
Current Load Balance: none  
Logical Unit Group ID: NA  
Auto Failback: on  
Auto Probing: NA
```

Paths:

Initiator Port Name:	10000000c93840a1
Target Port Name:	23000050ccf00e4a
Override Path:	NA
Path State:	OK
Disabled:	no

Initiator Port Name:	10000000c93840a0
Target Port Name:	21000050ccf00e4a
Override Path:	NA
Path State:	OK
Disabled:	no



#### Target Ports:

```
Name:                23000050ccf00e4a
Relative ID:         0

Name:                21000050ccf00e4a
Relative ID:         0
```

6) Observe StorView and start “Logical Drive Statistics.” From here you can observe I/O running through each controller, and even more specifically, each Logical Drive.

### Failover/Failback

In the event of a failover situation, repair the cause of the failover, ensure the original paths are enabled then disable the current path, which will fail the path back to its original configuration.

```
root@glendry # mpathadm disable path -i 10000000c93840a0 -t 21000050ccf00e4a -l /dev/rdisk/c7t60050C-
C000F00E4A000000000000022Ad0s2
```

### SAN Mapping Information

```
root@glendry # luxadm -e port
/devices/pci@1,0/pci10df,fa00@4/fp@0,0:devctl          CONNECTED
/devices/pci@1,0/pci10df,fa00@4,1/fp@0,0:devctl      CONNECTED
```

```
root@glendry # luxadm -e dump_map /devices/pci@1,0/pci10df,fa00@4/fp@0,0:devctl
Pos  Port_ID  Hard_Addr  Port WWN          Node WWN          Type
0    e1       0          21000050ccf00e4a  20000050ccf00e4a  0x0 (Disk device)
1    1        0          10000000c93840a0  20000000c93840a0  0x1f (Unknown Type,Host
                               Bus Adapter)
```

```
root@glendry # luxadm -e dump_map /devices/pci@1,0/pci10df,fa00@4,1/fp@0,0:devctl
Pos  Port_ID  Hard_Addr  Port WWN          Node WWN          Type
0    e1       0          23000050ccf00e4a  20000050ccf00e4a  0x0 (Disk device)
1    1        0          10000000c93840a1  20000000c93840a1  0x1f (Unknown Type,Host
                               Bus Adapter)
```

**What this shows** — The HBA port WWN for the entry with “(Unknown Type, Host Bus Adapter)” is the HBA you just queried with the command. This is also the value you would use in a switch to map the HBA. The other entries show the controllers seen by that HBA/port. This is an example of direct attach.

## Performance

None vs. Round Robin — To illustrate the performance impact between the load-balance settings of “round-robin” and “none” in the /kernel/drv/scsi\_vhci.conf file:

A test was run consisting of a cp of a 2.4GB file from the root drive to /sumo

-3 LUNs w/ 32KB chunk/416KB stripe/13 + 2

7 runs done for each setting.

load-balance=none:

```
real      53.2
real      50.0
real      46.5
real      45.9
real      46.0
real      45.9
real      50.1
```

average=48.23 sec. or 50.96MB/second

load-balance=round-robin:

```
real      1:13.8
real      1:20.0
real      1:25.3
real      1:12.3
real      1:17.2
real      1:31.2
real      1:11.7
```

average=1:18.80 min or 31.19MB/second (a %38.8 decrease.)

## About Xyratex

Xyratex is the ultimate partner to the storage industry. We are a leading provider of enterprise-class data storage subsystems and storage infrastructure manufacturing equipment & automation solutions. Working with over 50 A-list companies, Xyratex ships over 14% of the world's external storage capacity, and 75% of all 3.5" drives are processed using Xyratex test systems. With unmatched expertise and a history of innovation and technological excellence, Xyratex delivers products which are high-performance, energy-efficient and extremely reliable.

For more information, please visit [www.xyratex.com](http://www.xyratex.com)

### Xyratex Headquarters

Langstone Road  
Havant  
Hampshire PO9 1SA  
United Kingdom

### UK HQ

T +44 (0)23 9249 6000  
F +44 (0)23 9245 3654

[www.xyratex.com](http://www.xyratex.com)

### Principal US Office

2031 Concourse Drive  
San Jose, CA 95131  
USA

### USA Sales & Support

T +1 877 997 2839  
T +1 877 XYRATEX



ISO 14001: 2004 Cert. No. EMS91560

©2008 Xyratex (The trading name of Xyratex Technology Limited). Registered in England & Wales. Company no: 03134912. Registered Office: Langstone Road, Havant, Hampshire PO9 1SA, England. The information given in this brochure is for marketing purposes and is not intended to be a specification nor to provide the basis for a warranty. The products and their details are subject to change. For a detailed specification or if you need to meet a specific requirement please contact Xyratex: [www.xyratex.com](http://www.xyratex.com).

x y r a t e x