



Linux on IBM @server pSeries SAN Overview for Customers

February 07, 2005

Abstract

This paper is intended to introduce IBM @server® pSeries® customers, IBM Business Partners, sales, marketing, and technical teams to pSeries Storage Area Network (SAN) with Linux. IBM's plans as put forth in this document are subject to change without notice.

References and Prerequisite Reading

For an overview of Linux on pSeries, Emulex's Linux Fibre Channel device driver, and pSeries SAN solutions, the following background reading is recommended:

- The Linux for pSeries external Web site at <http://www.ibm.com/eserver/pSeries/linux>
- The IBM external Linux Web site at <http://www.ibm.com/linux>
- The IBM Linux Services Web site at <http://www.ibm.com/linux/services>
- The IBM SAN Storage Web site at <http://www.storage.ibm.com/ibmsan/>
- The IBM @server pSeries Server Web site at http://publib16.boulder.ibm.com/pSeries/en_US/infocenter/base/
- The Emulex FC MultiPulse Device Driver version 4.30o_M and 4.31j_M for IBM pSeries OEM-ed adapter (9802 and 9002L) at <http://www.emulex.com/ts/docoem/framibm.htm>
- Subsystem Device Driver (SDD) References
 - IBM Subsystem Device Driver User's Guide
 - Fault Tolerant Storage Multipathing and Clustering Solutions for Open Systems for the IBM TotalStorage® Enterprise Storage Server™ (ESS) is SG24-6295
 - IBM TotalStorage Enterprise Storage Server Implementing ESS Copy Services in Open Environments is SG24-5757
- FASTT References
 - IBM TotalStorage FASTT Storage Manager Concepts Guide, GC26-7560
ftp://ftp.software.ibm.com/pc/pccbbs/pc_servers/d5bajmst.pdf
 - IBM TotalStorage FASTT Storage Manager Copy Services User's Guide, GC26-7561
ftp://ftp.software.ibm.com/pc/pccbbs/pc_servers/fcs0mst.pdf
 - For hardware reference manuals, check the documentation link from the FASTT Technical support site at <http://www.ibm.com/servers/storage/support/fastt/>
 - IBM TotalStorage: FASTT Best Practices Guide
<http://publib-b.boulder.ibm.com/Redbooks.nsf/RedbookAbstracts/redp3690.html>

Overview of Linux on pSeries and Storage Area Network (SAN)

The SAN hardware supported by pSeries servers running the Linux operating system is the same SAN hardware that is supported by pSeries servers running AIX 5L™. The one exception is the IBM 6227 Fibre Channel adapter, as that adapter had been discontinued since the implementation of this document.

There may be differences between AIX 5L SANs and Linux on pSeries SANs. This document is designed to explain how to implement a Linux on pSeries SAN.

If a function is not explicitly stated to be supported in this document, it should not be assumed that it is supported; even if it is supported in an AIX 5L SAN. The FAQ section below tries to address the major differences between the two SANs. However, the FAQ may not cover all the differences between the two.

Overview of IBM pSeries Servers

This document covers all pSeries servers that support the Linux operating system and the 6228 and/or the 6239 Fibre Channel adapters. For server support, please refer to the pSeries documentation found at: http://publib16.boulder.ibm.com/pSeries/en_US/infocenter/base/

Certain features of pSeries servers that are supported under AIX 5L are not supported under Linux. Hot-swapping of a Fibre Channel adapter is one such example. The Linux 2.4.x kernel does not have the capability to power off/on a PCI slot and integrate newly added hardware as dynamically as AIX 5L.

The term “host”, “server”, and “system” have been used in this document and can be interchanged. The meaning of these words in this document is a single operating system instance. This could be a standalone server, or a single logical partition (LPAR) from a pSeries, which is hosting multiple LPARs. Two independent LPARs on the same physical hardware are considered two hosts (servers, systems) and should be treated as such.

Applicable SAN Environments for IBM’s pSeries Servers Running Linux

1. The Emulex adapter and driver are used in IBM @server iSeries™/pSeries with 64-bit PowerPC Architecture™ only. There is no support for older IBM RS/6000® or 32-bit PowerPC® machines.
2. Currently, the only supported Linux version is SuSE Linux Enterprise Server (SLES) 8 with SP3 applied.
3. Binary kernel versions:
 - a. 2.4.19-186 and later for Enterprise Storage Subsystem (ESS)
 - b. 2.4.21-138 and later for Fibre Array Storage Technology (FAStT)
4. For adapter placement and limitations please refer to the *PCI Adapter Placement Guide* (SA38-0538).
5. The Emulex device driver should be used with either Host Bus Adapters (HBA) IBM FC6228 or FC6239. The following versions of device drivers are supported, depending on the type of storage array in use:
 - a. v4.30o_M (ESS)
 - b. v4.31j_M (and later) (FAStT)
6. Disk subsystems
 - a. ESS Model 800 and F20
 - i. (R9e.8b030117) LIC EC ID 2.1.1.101 or above
OR
 - ii. (R9p.8b030606a) LIC EC ID 2.2.0.459 or above

- b. FAStT Models 100, 200, 500, 600, 600 Turbo, 700, and 900 with IBM FAStT Storage Manager v8.3, v8.4, or v9.1.
7. Currently, there is no FAStT Storage Manager Client Software available for pSeries host running Linux. To configure the FAStT disk subsystems, an AIX 5L, Windows® or PC-based Linux server is necessary.
8. SAN Topologies
 - a. Direct connect, via FC Arbitrated loop mode only, to ESS and FAStT.
 - b. Switch (Fabric) connect to ESS and FAStT.
 - c. No FC Arbitrated loop via hub connection is supported.
 - d. No direct connection with multiple paths (LUNs seen by multiple direct connections) is supported on ESS.
9. What things can NOT interoperate in the SAN environment?
 - a. No IP support
 - b. No ESS support on HBAs configured for MultiPulse.
10. Device sharing
 - a. A device/LUN can be shared with other HBAs on the same host as long as some type of multipath control is in place and an application is not using disk reservations.
 - b. A device/LUN can be shared with different hosts as long as all those hosts are pSeries hosts with Linux and there are no reservations used.
 - c. A device/LUN can not be shared with different hosts in a heterogeneous environment (one or more servers are NOT pSeries running Linux).

Emulex's Linux on pSeries MultiPulse Driver Capabilities

Emulex's Linux on pSeries MultiPulse driver provides single path and multipath capabilities for the FAStT disk subsystem, when configured accordingly. It also provides only single path capabilities for the IBM Enterprise Storage Server (ESS). Multipathing capability for the ESS is handled by the ESS's multipath SDD software. The Emulex MultiPulse driver does not provide multipath capability for the ESS.

Below is a list of capabilities for the Emulex MultiPulse driver in a pSeries host running the Linux Operating System. Some of these capabilities are restricted by entities outside of the driver, like the kernel or device restrictions.

- For the DS4000 series, up to 32 LUNs per FAStT per HBA, with a limit of 256 per server in a dual-path configuration
- Up to a queue depth of eight
- FAStT RDAC support only
- No support for FAStT Auto Volume Transfer (AVT) mode
- Path Failover Capability for FAStT disk subsystems, using versions 8.3, 8.4, or 9.10 of the IBM FAStT Storage Manager application.
- Multipathing capabilities via Switch attach only with ESS.

- Manual path recovery with ESS (automatic path recovery not yet supported)
- Up to 16 Emulex HBAs per host. pSeries server limitations may restrict this further.
- Both ESS and FAStT subsystems are supported on a single host, provided each subsystem is accessed via it's own HBA (FAStT and ESS are not accessed via the same HBA)
- IBM OEM-ed HBAs are the only supported Fibre Channel adapter. Emulex Fibre Channel adapters that are not OEM-ed by IBM will not work with the Emulex MultiPulse driver.
- Dynamic FAStT RAID-type changes are supported (e.g. converting a RAID-0 array to a RAID-5, or changing the RAID segment size while I/O is active).
- Dynamic volume expansions are supported on the storage. However, the new space may not be available for use dynamically, and will require reconfiguration of the devices and filesystems to recognize the add capacity of the volume.
- Dynamic LUN/SAN changes, which are not listed above, are not supported.
- I/O must be halted prior to making the following FAStT changes:
 - host type (i.e. AIX 5L to LNXCL)
 - primary and secondary path preference swapping

Different Driver Versions

Driver packaged with the Linux distribution (just load it and run):

- SLES8 GA: The Emulex Driver version 4.21f comes with the Linux distribution in SLES8 GA. This driver does not support MultiPulse (multi-path) devices.
- SLES8 SP1: The Emulex Driver version 4.22d comes with the Linux distribution in SLES8 Service Pack 1. This driver does not support MultiPulse (multi-path) devices.
- SLES8 SP3: The Emulex Driver in SP3 is the same driver version as SP1 (4.22d).
- MultiPulse: The Emulex Driver version 4.30*_M is MultiPulse enabled (provides multi-path functionality) and will fix three critical problems for the older drivers such as:
 1. Kernel panic when load or unloading the driver in race condition.
 2. Performance improvement from 64KB to 128KB SCSI buffer size.
 3. MultiPulse multi-path functionality.

The latest Emulex FC MultiPulse Device Driver (version 4.30o_M for ESS, 4.31j_M (and later) for FAStT) for IBM pSeries OEM-ed adapter (FC6239 =LP9802 and FC6228=LP9002L) can be downloaded from <http://www.emulex.com/ts/docoem/framibm.htm>.

Before Installing Emulex's Linux FC Adapter Driver

1. Storage: allocate Volumes/LUNs to each HBA on the pSeries host running Linux.
2. Switch: properly zone storage devices and the pSeries host running Linux; DO NOT allow any storage that is not configured for the pSeries host running Linux in its zone.
3. Server: it is recommended that the cable not be connected between the host and storage device unless there is a proper configuration created for the storage, including a matching

persistent binding configuration between the storage controllers and server HBAs compiled into the associated MultiPulse Emulex FC driver. You can load the Non-MultiPulse Emulex FC driver to verify the WWN and note the automatically driver assigned controller and target identifiers (look at “/proc/scsi/lpfc/*” files) to prepare for binding the intended MultiPulse or Non-MultiPulse storage configuration.

Adding a new FC Adapter

1. You MUST shutdown the pSeries host running Linux to install new Emulex FC Adapter. Linux does not have support for powering off an individual PCI slot.
2. After installing the new Emulex adapter, you MUST verify the MultiPulse configuration and re-compile the driver before loading the MultiPulse driver.

Quick Installation

There is no quick method for installation. Please follow the installation instructions below.

Installing and Configuring the Emulex FC MultiPulse Driver

The directions below must be followed in the exact order to guarantee the driver is built, installed, and loaded correctly. Whenever “/etc/lpfc.conf” is modified manually, or updated by the use of “hbacmd” or “hbanyware”, the lpfcdd driver must be re-compiled with the resulting modified “/etc/lpfc.conf” file before the bindings become effective.

1. Create temporary directory to build Emulex Driver and copy the driver you want to use into that directory (e.g. /emulex)
2. Install the SLES8 SP1 or SP3 kernel source code from SuSE distribution and also install the PowerPC GCC 64-bit compiler on the target pSeries host running Linux.

For instructions on how to install the kernel, see the Frequently Asked Questions section. Look for Question: How do I install the kernel and build driver for SuSE SLES8 SP3 for ppc64.

Note: There are 2 questions, one for SP1, and one for SP3. Choose the section that applies to the kernel level you are currently running. For the 4.31j_M (and later) driver, you should be running SLES8 SP3.

To install the PowerPC GCC 64-bit compiler and build environment, you can run YaST2 and choose the "64-bit Cross Compile Environment". Also choose any sections relevant to the build environment.

To verify you have the PowerPC cross compiler, run:

```
# rpm -qa | grep cross
```

You should see output similar to this with perhaps different version numbers:

```
cross-ppc64-gdb-5.2-115  
cross-ppc64-binutils-2.13.90.0.4-84
```

```
cross-ppc64-glibc-2.2.5-132
cross-ppc64-libs_and_headers-8.1-157
cross-ppc64-gcc-3.2.2-46
```

If the cross compiler is installed, it should reside under /opt/cross.

3. Ensure you have the correct kernel source and build environment for the Emulex MultiPulse Driver as detailed in FAQ.
4. Build the kernel dependencies. The instructions are given below if you are running SLES8 SP3 kernel. They are also duplicated in the FAQ section. If you are running SLES8 SP1, instructions are in the FAQ.

To build SLES8 SP3 kernel dependencies:

```
# cd /usr/src/linux
# make mrproper
# cp /boot/*.version.h /usr/src/linux/include/linux/version.h
# cp arch/ppc64/defconfig.pseries64 .config
# chmod 755 make_ppc64.sh
# sh make_ppc64.sh oldconfig dep
```

NOTE: Use the default answers when asked questions from oldconfig.

NOTE: If include/linux/modversions.h was not created:

- 1) Verify CONFIG_MODVERSIONS is set in your /usr/src/linux/.config file. CONFIG_MODVERSIONS=y is what you should see. If it is not set, set it.
- 2) Either way, re-run sh make_ppc64.sh oldconfig dep.

5. Fetch the Emulex Multipulse driver package. There are 2 flavors of files: a *.tar file, and a *.rpm. Choose whichever you are most comfortable with. If you choose the tar file, un-tar the driver in a temporary Emulex directory (e.g. /emulex), using “tar -xvf lpfc-4.31j_M-1_gcc3x.ppc64.tar”. If you choose the rpm, issue “rpm -ivh lpfc-4.31j_M-1_gcc3x.ppc64.rpm” (it will unpack into /usr/src/linux/drivers/scsi/lpfc/SourceBuild; resolve any conflicts with earlier Emulex drivers before proceeding; in addition, there may be naming issues for certain files...see the FAQ).

6. Install and build the Emulex MultiPulse driver as with no MultiPulse devices (using the default lpfc.conf) configured as follows:

Change directory to where the Emulex Driver source code resides

```
# chmod 755 Install.sh
# sh Install.sh
# make build
# make install
```

7. Load the driver, making sure that you have the storage and switch configured correctly. Load the Emulex driver with the command “modprobe lpfcdd”

- a. Confirm that the driver loads correctly with the command “lsmod” and look for the driver name “lpfcdd”

- b. NOW, you can connect the fibre cable to HBA and see if the Linux on pSeries is able to detect the switch and storage devices with the following commands:

```
# ls /proc/scsi/lpfc/          * check number of HBA installed
```

```
# cat /proc/scsi/lpfc/1      * check status of 1st HBA
# cat /proc/scsi/lpfc/2      * check status of 2nd HBA
# cat /proc/scsi/lpfc/3      * check status of 3rd HBA,...
```

- c. If the driver is correctly built and loaded, you should be able to see the WWPN and WWNN of all connected storage devices listed in the /proc/scsi/lpfc/* files. Verify fabric and loop connections. (Make sure that if you are using fabric zoning that the devices are where they are supposed to be.)
- d. Now double-check to see if all the storage subsystem with its WWNNs and WWPNs are seen by the HBA and their devices are configured for the pSeries host running Linux. If they are correct, you can unload and reload the driver quickly and ready to use it.

```
# rmmod lpfcdd              * Unload the Emulex driver
# modprobe lpfcdd           * Load the Emulex driver
```

8. Now install and build the Emulex MultiPulse driver with multiple paths to the storage devices. **IMPORTANT: ALWAYS CONFIGURE PERSISTENT BINDINGS BY WWPN AND NOT BY WWNN NOR BY DOMAIN ID AS THE /etc/lpfc.conf DOCUMENTS AS OTHER BINDING METHODS (WWNN and DID) do not work).**

- a. After you are done with previous steps to install the Emulex driver as non-MultiPulse devices, you should have obtained all the information needed to rebuild the Emulex driver with MultiPulse device support.
- b. Note all the WWNN and WWPN associated with the HBAs
- c. Rebuild the Emulex driver with `make build` then `make install` commands as outlined earlier from the local Emulex driver source code directory. There are 3 different ways to configure the driver, listed below. Choose one. *However, IBM recommends running hbacmd.*
 - i. Use command line interface program `/usr/sbin/hbanyware/hbacmd mpMenu <AdapterWWPN>` to create the new `/etc/lpfc.conf` (see section on `hbacmd` utility under “Driver Customization” section below);
 - ii. Edit the `/etc/lpfc.conf` to bind the HBA to the Storage Devices with its WWPN (a configuration may be restored by reinstalling a saved copy of “`lpfc.conf`”);
 - iii. Use GUI interface program “`/usr/sbin/hbanyware/hbanyware`” to create the new `/etc/lpfc.conf`. (see section on `hbanyware` utility). Java must be installed to run the GUI interface program.
- d. Load the Emulex driver via `modprobe lpfcdd` command.

Driver Customization

Setting up consistent device numbering (implications of changing configurations and rebooting)

- Consistent device numbering only for Emulex binding parameters. The device instances may not be consistent after booting or reconfiguration if SAN or storage changes have been made. However, there are free Linux utilities that may help provide consistent device numbering or consistent device instances, such as `scsimap`.

All ESS and FAStT devices may have persistent bindings configured for them. Persistent bindings help to ensure that specific WorldWide PortNames are always given the same SCSI target id by the lpfcdd driver, which provides consistency and easier storage management.

It is highly recommended that you use persistent bindings. To configure the /etc/lpfc.conf file to manage only those storage targets with persistent bindings, edit the “lpfc_automap” variable in /etc/lpfc.conf to equal 0 (i.e. “lpfc_automap=0”).

Configuration of persistent bindings should be done before configuring MultiPulse routes and can be done one of 2 ways:

1. Modify the /etc/lpfc.conf file according to the instructions found therein;
2. Use hbacmd utility’s “Persistent Bindings Menu”.

Choose “Bind Automapped Targets” which instructs the driver to bind the targets to adapter and target instances that it prefers to select. Each adapter must be automapped this way. Also, you must bind all targets by WWPN, or “By Port Name”. (The other methods will not work, which includes binding by World Wide Node Name or by Domain ID.)

You can then choose “Delete Binding(s)” and remove the storage targets you do not want. By selectively removing bindings, and setting the “lpfc_automap” variable in /etc/lpfc.conf to equal 0, you can control which storage devices the driver will manage. In other words, you can control which targets your pSeries host running Linux will “see” each time the driver is loaded, including the boot-time scan. This control over driver load times can reduce boot times.

Persistent Bindings Notes:

- All bindings must be by WWPN.
- When manually editing /etc/lpfc.conf for bindings, first compile and load lpfcdd with the default lpfc.conf file, and use the adapter and target instances found in the /proc/scsi/lpfc/* files. During testing it has been observed that using target ids other than those “preferred” by the lpfcdd driver will not work.
- When using persistent bindings, the /etc/lpfc.conf file must be manually edited and the variables “lpfc_automap” and “lpfc_scandown” set equal to zero (“0”). With automap set to 0, only those devices with defined persistent bindings will be found, so make this edit *after* the initial discovery.
- Remember that changes to the persistent binding configuration (as denoted in the lpfc.conf file) require recompilation of the lpfcdd driver. Otherwise, the changes will have no effect.

Configuration of MultiPulse routes can be done in one of 3 different ways.

1. Manually edit the /etc/lpfc.conf file according to the instructions found therein.
2. Use the hbacmd utility’s MultiPulse menu, using the command:

```
# /usr/sbin/hbanyware/hbacmd mpMenu <AdapterWWPN>
```

The WWPN and lpfc# of each adapter can be found its specific /proc/scsi/lpfc/* adapter configuration file.

Within the MultiPulse menu choose “1. Create Route”

select the LUN for which you want multi-path support, for example “1. lpfc0t1d0”

select the path to include in the route, for example “1. lpfc0t1d0”

select the second path to include in this route, for example “2. lpfc2t4d0”

select

choose “y” to enable AutoFailBack; (Unless using FlashCopy: see section on FAST copy services.)

choose “n” for InquiryHeartBeat, FailoverAbortTaskSet, FailoverLunReset, and FailoverTgtReset (Use of these features is not supported yet.)

select “1. Simple Failover” for the control type

select the path you want to be the “primary path”, or preferred path for I/O.

For example 2, “lpfc0t1d0” will cause I/O to run to that device when driver is loaded, regardless of how you have configured the FAST Preferred Controller Ownership for that volume using the SMclient graphical user interface.

(Hint: If you alternate the paths here you will essentially load-balance your I/O.)

answer “Y” to “Are you sure you want to create this route?”

Actually, you can make short work of this by recording this string of responses in a separate text file, then using a shell command to redirect that file into the hbacmd utility. For example, entering the string

```
“1 1 1 2 0 y n n n n 1 1 y”
```

sequentially after invoking hbacmd would result in a route being created as specified above. Repeating this line in the text file provides a “one touch” means of creating a number of routes: ten such lines would result in up to 10 routes created each time the following command was issued.

```
# /usr/sbin/hbanyware/hbacmd mpMenu <AdapterWWPN> <textfile>
```

See how the file “addroute” is used in Section 5 of Appendix A for a more complete example.

Be careful using this method. Its best to carefully check the sequence before using this method to configure a large system. In addition, there is an issue with the hbacmd utility that may cause problems if the lower ordered of two lpfc's is not selected as the primary path. In that case, when the script begins to cross boundaries from a set of luns on one FAST to those configured on another FAST, LUN 0 may get omitted. After running the shell command, always verify that the mappings are listed correctly in the lpfc.conf file, in the section labeled “BEGIN: MultiPulse managed entries”.

3. Use the hbanyware graphical user interface, starting it using the command below, making sure that your DISPLAY environment variable is correctly set:

```
# /usr/sbin/hbanyware/hbanyware
```

In resulting GUI:

Find HBA WWN in "Discovered Elements" list

click "+" symbol next to storage icon with WWNN you want to manage, then

click on "LUN 0" -> "MultiPulse" tab

In the MultiPulse tab section, click on the "Avail. Path ID:" for the path you want to be the primary path. This should be the path to the controller which is the preferred owner of the volume mapped to this LUN instance, and then click "Create MP Route" check box.

click "Apply Checked Item" button to finish configuring this route.

Verify the route is configured by clicking on the "+" symbol next to the storage icon which represents the alternate controller which also owns this LUN, and check the "MultiPulse" tab. This route should be represented.

Do this for each LUN the controller owns, making sure that the first path you add by clicking on "Avail. Path ID:" is the desired primary path.

Commonly customized parameters found in "lpfc.conf"

```
lpfc_fcp_bind_WWNN[]=... // array of storage WorldWideNodeNames
                          // and their bindings
lpfc_fcp_bind_WWPN[]=... // array of storage WorldWidePortNames
                          // and their bindings
lpfc_fcp_DID[]=...       // array of storage DestinationIdentifiers
                          // and their bindings (note that these 3
                          // binding methods are mutually exclusive)
lpfc_automap=0;          // Set to 0 for persistent binding
lpfc_scandown=0;         // Set to 0 for persistent binding
lpfc_multipulse_entries=<n>; // "n" MUST match number of entries in
                          // "lpfc_multipulse_fcp[]"
lpfc_multipulse_fcp[]    // array of LUNs with their failback paths
lpfc_linkdown_tmo=1;     // Link timeout value; set in lpfc.conf.c
lpfc_nodev_tmo=1;        // "No Device" timeout value; set in
                          // lpfc.conf.c
```

Also, the amount of time required to boot Linux can be reduced by adding the following line to the boot loader configuration file (e.g. /etc/lilo.conf or /etc/yaboot.conf):

```
append = "max_scsi_lun=16" // limit SCSI scan to 16 devices
```

For further advanced configuration, please refer to the Emulex driver documentation or see Appendix A for some examples.

Remove or Reinstall the Driver

To remove (unload) the driver, issue "rmmod lpfcdd"

To reload the driver, issue "modprobe lpfcdd"

To reinstall the driver, follow the section above on how to build and install the Emulex driver.

Setting up Multiple Paths for FAStT Storage

The Emulex MultiPulse driver supports multipath for FAStT storage subsystems only. There are different models of FAStT controllers, and all of them are supported (FAStT100, 200, 500, 600, 600 Turbo, 700, and 900). Below is a list of considerations:

- For multi-path support, the IBM FAStT Storage Manager must be at version level 8.3, 8.4, or 9.10.
- The Host Type for FAStT Linux on pSeries connections is LNXCL.
- Connect and configure multiple paths to a FAStT using the IBM FAStT Storage Manager Client.
- Up to two paths to a FAStT are supported, no more than one per controller.
- If you use only FAStT and Multipath Devices, make sure to set `lpfc_automap=0` and `lpfc_scandown=0`.
- Do not configure multipath for ESS devices using the Emulex MultiPulse driver. Make sure to set `lpfc_automap=2` and `lpfc_scandown=0`.

Path Failures

ESS Environments:

There is no multipath support for ESS with the Emulex driver. The ESS device will be treated as single path device. When the path fails, the Emulex driver will report an I/O error to the upper SCSI layer.

FAStT Environments:

If a FAStT volume is configured as a MultiPulse device, and the primary path fails, the MultiPulse driver will automatically failover to the secondary path.

If all paths are failing, the Emulex driver will report an I/O error to the upper SCSI layer.

Recovery from failed paths and offline paths is the same. There is no difference between failed paths and offline paths.

Installing and Booting from a SAN Device

SAN Installation and booting has not been fully tested and is not currently supported by pSeries servers running Linux.

Copy Services

ESS

There are no limitations on ESS copy services beyond the ESS copy services limitations defined for i386 Linux attached servers. Refer to the ESS copy services document [SG245757](#) for information on ESS copy services in a Linux Environment

FAStT

Other than the following notes on Path Failover/Failback with FlashCopy there are no limitations on FAStT copy services beyond the FAStT copy services limitations defined for

i386 Linux attached servers. Refer to the *FAStT Storage Manager Copy Services User's Guide* for information on FAStT copy services in a Linux Environment

There are two types of failover/failback scenarios for FAStT FlashCopy:

1. Failover/Failback of an entire path

The Failover/Failback of a path is straightforward and works as expected. See the later section in this document titled "Path Failures" for more information on this type of Path Failover/Failback.

2. Failover/Failback of a LUN (LUN ownership)

If LUN Ownership is changed for a FlashCopy Base volume the FlashCopy Logical volume LUN ownership will also change. In order to change the ownership back to the original configuration all I/O to the FlashCopy Logical Volume must be stopped while you manually reconfigure the LUN ownership and put the paths to the FlashCopy Logical Volume online.

Auto Failback should not be used with paths that are involved with FlashCopy Logical Volumes.

Helpful Commands and Tools

```
cat /proc/scsi/scsi           // show SCSI devices
ls /proc/scsi/lpfc/          // show Emulex adapter number
cat /proc/scsi/lpfc/1        // show detailed information for Emulex adapter 1
cat /proc/scsi/lpfc/2        // show detailed information for Emulex adapter 2
fdisk                        // utility to show, create, format disks
yast2                        // SLES8 system management GUI
lputil                       // Emulex command line diagnostic utility
hbacmd                       // Emulex command line configuration tool
hbanyware                    // Emulex GUI configuration tool
```

Frequently Asked Questions

Q: Why am I missing some, or all, of the LUNs/devices on my storage subsystem?

A: If you do not see any working LUNs/devices, you should check the following:

1) You may not have installed and loaded the Emulex driver. Check if the Emulex driver is loaded by using “lsmod” to list all the driver modules that were loaded. If you see the lpfcdd entry on the list, the Emulex driver was loaded. Otherwise, you can use “modprobe lpfcdd” command to load the driver.

2) If you have loaded the driver properly and it is shown as in “lsmod” list, check the following SAN configuration issues:

- fibre cable is properly connected to Emulex HBA, to Switch, and to storage subsystem.
- switch is properly zoned to recognize the Emulex HBA WWNN and storage WWNN
- storage subsystem has configured LUN/devices and associated them with the Emulex HBA WWNN/WWPN
- For **ESS** subsystem: Host type is set to i/Linux on pSeries for the LUNs
- For **FAST** subsystem: Host type is set to “LNXCL” (Linux Cluster).

3) LUN 0 must be defined in order for device discovery to work properly. Configurations without LUN 0 configured are unsupported. If you do not have LUN0 defined, you need to specify lpfc_missing_lun0=1 in lpfc.conf.c before you build and load the Emulex driver (you can also load the parameter at the time of loading the driver as “modprobe lpfcdd lpfc_missing_lun0=1”). However, if LUN0 is added later, the subsequent rediscovery (via “rmmod” / “modprobe”) will fail.

4) If you have non-consecutive LUNs defined (i.e. LUN0, LUN2, LUN3 ... and do not have LUN1 defined), you need to specify lpfc_lun_skip=1 in lpfc.conf.c before you build and load the Emulex driver. You can also load the parameter at the time of loading the driver as “modprobe lpfcdd lpfc_lun_skip=1”.

5) Check to see if the driver can detect the switch and ESS storage subsystem properly by using these following commands.

```
# ls /proc/scsi/lpfc           // check if driver sees any Emulex adapter
# cat /proc/scsi/lpfc/(n)     // list if driver sees any target devices
# cat /proc/scsi/scsi         // lists all scsi devices
```

- If it shows that the link is down, it means there is cable problem between Emulex HBA and switch, or bad cable. “lpfc” should show the link as fabric or loop.
- If it shows that the link is up but did not show any target WWNN/WWPN, it means there is cable problem between switch and ESS storage subsystem; or bad cable; or incorrect Switch zone setting.
- If it shows that the link is up and target WWNN/WWPN, it means there is incorrect configuration of host type; or Emulex HBA WWNN associated with LUNs/devices; or missing LUN0 defined and the lpfc_missing_lun0 was set to 0. Another possibility is the mapping of HBA WWPN on the FAST is incorrectly set.

Q: Why is there such a long delay when loading the Emulex Driver?

A: If there is a device connection seen by the pSeries host running Linux, which does not have any LUNs configured, the Emulex driver may experience long time outs as it tries to find a device. To clear up this issue, make sure each connection seen by the HBA has at least one device configured.

Q: Why do I see more devices than are configured?

A: If there is a device connection seen by the pSeries host running Linux, which does not have any LUNs configured, the Emulex driver may attempt to configure a LUN for each LUN ID. These are ghost devices and are not valid drives. To clear up this issue, make sure each connection seen by the HBA has at least one device configured.

Q: Why does the system panic when I try to unload the driver?

A: If there is a device connection seen by the pSeries host running Linux, which does not have any LUNs configured, the Emulex driver may cause a kernel panic. To clear up this issue, make sure each connection seen by the HBA has at least one device configured.

Q: How to I assign FASTT devices to a pSeries host running Linux?

A: Currently, there is no FASTT Storage Manager Client Software available for pSeries host running Linux. To configure the FASTT disk subsystems you will need to use an AIX 5L, Windows or i386 processor-based Linux server, and run a version of the FASTT Storage Manager client that is supported on that server.

Q: How do I configure MultiPulse devices (multipathing devices) for a pSeries host running Linux using the Emulex FC MultiPulse driver?

A: When you already have the FASTT configured with MultiPulse devices, you can load the Emulex FC MultiPulse driver as single path device driver. This single path driver will be able to detect the MultiPulse device as two independent, separate devices. Using “hbacmd” or “hbanyware”, you can configure these two independent devices to be a Multipath device. The “hbacmd” or “hbanyware” will create a new /etc/lpfc.conf with matching configuration for a MultiPulse device. You now can rebuild the new Emulex FC MultiPulse driver with the new /etc/lpfc.conf. After you build the driver with new /etc/lpfc.conf, you can unload the current single path driver and reload the new driver to use the device as MultiPulse device.

Q: Why is the MultiPulse driver not managing paths?

A: If the MultiPulse driver is not managing paths, the devices were probably not configured for MultiPulse operation. Refer to the question above on how to set up MultiPulse operation.

Q: Can data be shared across adapters and/or servers?

A:

1. A device/LUN can be shared with other HBAs on the same host as long as some type of multipath control is in place and disk reservations are not used.
2. A device/LUN can be shared with different hosts in a homogeneous environment (ALL servers are Linux on pSeries based) as long as there are no disk reservations.
3. A device/LUN can not be shared with different hosts in a heterogeneous environment (one or more servers is NOT Linux on pSeries based).

Q: Which Fibre Channel Host Bus Adapters are supported?

A: This release supports only the IBM FC6239 (LP9802) and IBM FC6228 (LP9002L) HBAs (the FC5716 (LP10000) has not been fully tested, and so is not supported).

Q: Where do I get the driver?

A: You can download the latest Emulex FC MultiPulse Device Driver for IBM pSeries OEM-ed adapters at <http://www.emulex.com/ts/docoem/framibm.htm>.

Q: How do I install the kernel and build the driver for SuSE SLES 8 SP1 for ppc64?

NOTE: If you are running SLES8 SP3, go to next question. This is for SLES8 SP1.

A: You should follow the instructions in the section above on how to build the driver. However, here is a quick list of the necessary operations:

Install the kernel source and ppc64 C++ compiler matching with current binary kernel version. To check for current kernel in use, you can use “uname -a” for active kernel in use; and use “rpm -qa | grep kernel” for current kernel binary and source installed on your system. To install a kernel source rpm, you can use rpm command “rpm -ivh kernel-source-2.4.19-SuSE-xxx.ppc.rpm”. To upgrade a kernel source rpm, you can use rpm command “rpm -Uvh kernel-source-2.4.19-SuSE-xxx.ppc.rpm”.

Build kernel dependency for driver build

```
# cd /usr/src/linux
# cp /boot/*.version /usr/src/linux/include/linux/version.h
# chmod 755 make_ppc64.sh
# sh make_ppc64.sh cloneconfig dep
```

Download the driver and un-tar it to local directory.

```
# tar -xvf lpfc-4.30o_M-1_gcc3x.ppc64.tar
```

Build driver from this local directory

```
# chmod 755 Install.sh
# sh Install.sh
# make build
# make install
```

Load the driver

```
# modprobe lpfcdd
```

NOTE: If include/linux/modversions.h was not created:

- 1) Verify CONFIG_MODVERSIONS is set in your .config file. CONFIG_MODVERSIONS=y is what you should see. If it is not set, set it.
- 2) Either way, re-run sh make_ppc64.sh oldconfig dep.

Q: How do I install the kernel and build the driver for SuSE SLES 8 SP3 for ppc64?

A: You should follow the instructions in the section above on how to build the driver. However, here is a quick list of the necessary operations:

Install the kernel source and ppc64 C++ compiler matching with current binary kernel version. To check for current kernel in use, you can use “uname -a” for active kernel in use; and use “rpm -qa | grep kernel” for current kernel binary and source installed on your system. To install a kernel source rpm, you can use rpm command “rpm -ivh kernel-source-2.4.21-xxx.ppc.rpm”. To upgrade a kernel source rpm, you can use rpm command “rpm -Uvh kernel-source-2.4.21-xxx.ppc.rpm”.

Build kernel dependency for driver build

```
# cd /usr/src/linux
# make mrproper
# cp /boot/*.version.h /usr/src/linux/include/linux/version.h
# cp arch/ppc64/defconfig.pseries64 .config
# chmod 755 make_ppc64.sh
# sh make_ppc64.sh oldconfig dep
```

Download the driver and un-tar it to local directory.

```
# tar -xvf lpfc-4.30o_M-1_gcc3x.ppc64.tar
```

Build driver from this local directory

```
# chmod 755 Install.sh
# sh Install.sh
# make build
# make install
```

Load the driver

```
# modprobe lpfcdd
```

NOTE: Take all of the default answers when asked questions from oldconfig.

NOTE: If include/linux/modversions.h was not created:

- 1) Verify CONFIG_MODVERSIONS is set in your .config file.
CONFIG_MODVERSIONS=y is what you should see. If it is not set, set it.
- 2) Either way, issue sh make_ppc64.sh oldconfig dep.

Q: I follow the instruction to build the driver but when I load the driver, it gets error messages: “kernel version mismatch” and the driver is not loaded. What should I do?

A: If you see the kernel version mismatch error, check the following:

- what is the active kernel in use by issue “uname -r”;
- what is the defined version in kernel source by issuing
#cat /usr/src/linux/include/linux/version.h
- If it does not match the “uname -r” version, then make sure that the kernel dependency was built correctly.

After you have matching kernel version from “uname -r” and version.h; you should be able to rebuild, install and load the driver correctly.

Q: When I load the driver via “modprobe lpfcdd”, it takes a long time and seems like it never loaded. What should I do?

A: When you don't have a storage device configured, but you already have connected all of the cables, the driver will take a long time trying to discover devices on the SAN network which are not there or unknown. It may also try to communicate with phantom or unknown devices and will not receive a response, which will cause timeout, and/or retry. You should disconnect the fibre cable at the Emulex HBA, then wait until the driver finishes loading. Then, follow the instructions from previous chapter on how to configure and install the driver.

Q: When I load the driver via “modprobe lpfcdd”, the system panics. What should I do?

A: You should make sure that you have a correct kernel installation. Do not use unknown kernel images or those built on a different system. If you have a good Linux kernel, make sure that you

have correctly installed the Emulex driver. If you have a good Emulex driver installed, make sure that you have the latest Emulex adapter firmware, and only use IBM OEM-ed adapters.

Q: When I unload the driver via “rmmod lpfcdd”, the system panics. What should I do?

A: You should make sure that the driver has been built with these parameter settings in lpfc.conf.c as lpfc_missing_lun0=1 and lpfc_lun_skip=1, especially when you have a storage subsystem connected to the SAN network but is not configured for a pSeries host running Linux.

Q: I have built everything correctly, but I still cannot get the adapter configured.

A: Make sure the adapter is an IBM OEM-ed adapter. The Emulex MultiPulse driver has been designed to work with IBM OEM-ed adapters only.

Q: Shouldn't the performance be better than I am seeing?

A: The Linux kernel has some limitations when it comes to SCSI based protocol devices. Mainly those drivers must be configured to operate on a single CPU only and have a buffer limitation of 128K. Tuning general SAN components, like queue depth, may also help performance.

Q: Can I hot swap a Fibre Channel adapter like I can on AIX 5L?

A: No. Linux does not support PCI slot power control, which is needed for hot swapping of PCI adapters.

Q: Are there any known issues (and fixes) and their workarounds with the Emulex driver?

A: Here is a list of the known problems with Emulex driver v4.30o_M driver:

1. (T082) FlashCopy does not work properly with autofailback. Needs to disable autofailback and manual intervention from user with hbacmd or hbanyware to move secondary copy back. (Fixed in v4.31h_M)
2. (T068) The Emulex 4.30o_M MultiPulse Driver only supports up to 16 Emulex HBAs per pSeries host running Linux.
3. (T080) If you build the MultiPulse driver with predefined MultiPulse devices in lpfc.conf and do not have matching storage network devices matching; you will not be able to load the driver correctly. You should build the driver with default lpfc.conf; then load this driver and use “hbacmd” or “hbanyware” to configure MultiPulse devices with this driver; after you have create a new lpfc.conf, you can build and load this new MultiPulse driver.
4. (T078) If you try to unload the driver and receive error: “<3>Attached usage count = -1”; you should rebuild the driver with lpfc_missing_lun0=1 and lpfc_lun_skip=1 in lpfc.conf.c; or load the driver as “modprobe lpfcdd lpfc_missing_lun0=1 lpfc_lun_skip=1”.
5. (T077) If you run application and receive kernel panic with “PCI_DMA: TCE allocation failure in create_tces_sg”; you should build the driver with lpfc_lun_queue_depth with smaller value in lpfc.conf.c, lpfc_lun_queue_depth=8 is recommended; or load the driver as “modprobe lpfcdd lpfc_lun_queue_depth=8”.
6. (T067) (T075) If you are using SuSE SLES8 SP1 kernel and experience system hung unresponsive but you still can ping to the system; you should build the driver with lpfc_lun_queue_depth with smaller value in lpfc.conf.c, lpfc_lun_queue_depth=8 is recommended; or load the driver as “modprobe lpfcdd lpfc_lun_queue_depth=8”. If you have more than 32 LUNs/devices configured for an Emulex HBA, you should only have less than 32 LUNs/devices in used concurrently at a time. It is best to have 32 or less LUNs/devices per Emulex HBA.
7. (T074) If you are unable to configure MultiPulse devices for FASTt storage subsystem with MultiPulse Emulex HBAs, make sure that you have IBM OEM-ed adapters. The Emulex MultiPulse driver only supports IBM OEM-ed adapters on pSeries.
8. (T072) (T073) If you experience an application, using MultiPulse devices, running but no active I/O data transfer, you may run into problem with switch port failing with constant link down/up or RSCN and does not failover to secondary path; you should check the switch port for errors. If you determine that the switch port is indeed having the problem, you can 1) disconnect the fibre cable from the bad port and move to different working port; 2) disconnect the fibre cable from the bad port; 3) disable the primary path via “hbacmd” or “hbanyware” utilities.

9. (T071) If you see lot of messages “!lpfc(x):0756:FPw:Local_Timeout” in /var/log/messages; you should check to make sure that the driver have lpfc_log_verbose set to 0 in lpfc.conf.c; or load the driver as “modprobe lpfcdd lpfc_log_verbose=0”.
10. (T069) If you load the driver and receive error “bad virtual addr” in /var/log/messages; you should make sure that you have enough free memory available in the system. You should check if there are too many processes running and you are running out of memory. Use “cat /proc/meminfo” to see how much memory available; you can also use “top” to see how your system is running.
11. (T060) If you experience a long delay after you have closed or killed the application using I/O devices via Emulex adapters; you should check to see if there are any error messages posted in /var/log/messages; these errors may cause long delay in cleaning up the memory and I/O request queues before the application is actually closed.
12. (T062) If you want to configure FAStT with host type Linux and RDAC enable but the default was AVT; you should configure the host type as LNXCL (Linux Cluster); or you can use script to enable RDAC for host type Linux.
13. (T041) If you are unable to build and install Emulex driver on SuSE 2.6 kernel or Red Hat; the current Emulex 4.30o_M driver only support SuSE 2.4 – SLES8 SP1 or SP3.
14. (T080) If you do not see any path or route after loading a MultiPulse driver; probably the MultiPulse driver does not match the current SAN configuration (i.e. the SAN configuration has changed, a LUN/device has been removed, a path is offline, a fibre cable has been disconnected. You can use “hbanyware” to reconfigure the MultiPulse devices or you can Unload the current driver, Reload the single path driver, and Rebuild a new MultiPulse device configuration and then load this new driver.
15. (T081) The system is panic when you try to unload the driver with 16 HBAs. This problem is still under investigation and may be fixed in a future release.
16. (T084) Swapping primary and secondary paths preferences while I/O active does not work. This problem is known and will be fixed in next release. The current solution is to stop all I/O before swapping primary and secondary paths. ([Fixed in v4.31h_M](#))
17. (T085) No MultiPulse devices entries available on primary HBA is bad. This problem is still under investigation and may be fixed in a future release. The current solution is to replace the bad HBA.
18. (T086) When failback occurs, one device or more MultiPulse devices were inaccessible because the primary paths for these MultiPulse devices were offline. This problem is known and under investigation. The current solution is to stop I/O to the failing devices and use “hbanyware” to enable the path that was offline. ([Fixed in v4.31h_M](#))
19. (T079) When failback, the system is panic with “kernel panic: NULL devp in flush_done”. The problem is fixed in 4.30o_M driver. However, if you are still experience the problem. You may have the FAStT and ESS sharing the same HBA in MultiPulse device configuration. You should move the ESS devices to different HBA.
20. (Bugzilla11320) - Reiserfs_panic: System panics when disk becomes unavailable. A panic will occur if a Raid 0 drive with a mounted reiserfs filesystem becomes unavailable when the filesystem tries access to the drive.

Q: The “Install.sh” script reports certain files not found (e.g. dfc)...how do I get around this?

A: Make sure that the files being reported as missing are present in the install directory. For example, a list of the directory may show both dfc32 and dfc64, but no dfc file. Create a symbolic link to the one that matches your platform type: “ln -s dfc64 dfc”. Do this for each file the script complains about. Then rerun the script.

Q: I’m seeing some IOs take an extended amount of time to complete...why?

A: If your system is under high I/O load and a storage system maintenance action is in process, you may see extended IO completion delays. If this behavior continues beyond completion of the maintenance action, please check the health of your system, or contact your service personnel for support.

Appendix A

Example of loading and building the Emulex FC Multipulse driver for FASSt devices

These steps should be done in the order listed below for a new installation.

1. Download and un-pack the Emulex FC Multipulse Driver
2. Build the Emulex FC Multipulse driver as a single path driver
3. Create persistent bindings for FASSt Multipulse devices
4. Single path configuration for FASSt Multipulse devices
5. Create entries for FASSt Multipulse devices in multi-path configurations
6. Build a Multipulse driver with FASSt multi-path devices
7. Final Step – load the new driver for multi-path configuration

(Note: the examples below show the results with lpfc-4.30o_M driver code; substitute the current released version in the commands. For example, the current version at the time of this writing is lpfc-4.31j_M.)

1. Download and un-pack the Emulex FC Multipulse Driver

```
# tar -xvf lpfc-4.30_M-1_gcc3x.ppc64.tar
lpfc-4.30o_M/
lpfc-4.30o_M/dfc
lpfc-4.30o_M/lpfcdriver
lpfc-4.30o_M/Makefile
lpfc-4.30o_M/fcLINUXfcp.c
lpfc-4.30o_M/README
lpfc-4.30o_M/fcLINUXlan.c
lpfc-4.30o_M/lpfc.conf.defs
lpfc-4.30o_M/Install.sh
lpfc-4.30o_M/Makefile.kernel
lpfc-4.30o_M/Makefile.module
lpfc-4.30o_M/Remove.sh
lpfc-4.30o_M/libemulexhbaapi.so
lpfc-4.30o_M/hbanyware/
lpfc-4.30o_M/hbanyware/librmapijni.so
lpfc-4.30o_M/hbanyware/rm.conf
lpfc-4.30o_M/hbanyware/hbacmd
lpfc-4.30o_M/hbanyware/hbanyware
lpfc-4.30o_M/hbanyware/emulexRMConfig
lpfc-4.30o_M/hbanyware/librmapi.so
lpfc-4.30o_M/hbanyware/HBAnyware.jar
lpfc-4.30o_M/lputil
lpfc-4.30o_M/lpfc.conf.c
lpfc-4.30o_M/lpfc_tar.spec
lpfc-4.30o_M/lpfc.conf
lpfc-4.30o_M/libdfc.a
```

```
lpfc-4.30o_M/include/  
lpfc-4.30o_M/include/fc.h  
lpfc-4.30o_M/include/dfc.h  
lpfc-4.30o_M/include/mpdevvid.h  
lpfc-4.30o_M/include/fc_hw.h  
lpfc-4.30o_M/include/fc_os.h  
lpfc-4.30o_M/include/fcdds.h  
lpfc-4.30o_M/include/fcmsg.h  
lpfc-4.30o_M/include/fcdiag.h  
lpfc-4.30o_M/include/mpapi.h  
lpfc-4.30o_M/include/mplib.h  
lpfc-4.30o_M/include/fc_crtn.h  
lpfc-4.30o_M/include/fc_ertn.h  
lpfc-4.30o_M/include/fcfgparm.h  
lpfc-4.30o_M/include/fcfgtgtm.h  
lpfc-4.30o_M/libHBAAPI.so
```

2. Build the Emulex FC Multipulse driver as single path driver

```
# cd lpfc-4.30o_M  
# sh Install.sh  
# make build >/dev/null 2>/dev/null  
# make install  
Installing under /lib/modules/2.4.21-111-pseries64/kernel/drivers  
# rmmod lpfcdd  
# modprobe lpfcdd  
Warning: loading /lib/modules/2.4.21-111-  
pseries64/kernel/drivers/scsi/lpfcdd.o will taint the kernel: no license  
See http://www.tux.org/lkml/#export-tainted for information about  
tainted modules  
Module lpfcdd loaded, with warnings
```

3. Create persistent bindings for FAStT Multipulse devices

```
# cat /proc/scsi/lpfc/2  
Emulex LightPulse LPFC Driver Version: 4.30o_M  
HBA: Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter  
SerialNum: 0000c9240da2  
Firmware Version: 3.91A1  
Hdw: 1002606d  
VendorId: 0xf90010df  
Portname: 10:00:00:00:c9:24:0d:a2 Nodename: 20:00:00:00:c9:24:0d:a2  
  
Link Up - Ready:  
  PortID 0x6c1213  
  Fabric  
  Current speed 2G  
  
lpfc0t00 DID 6c0813 WWPN 50:05:07:63:00:cc:0e:64 WWNN 50:05:07:63:00:c0:0e:64  
lpfc0t01 DID 6c0c13 WWPN 20:02:00:a0:b8:0c:cf:d2 WWNN 20:02:00:a0:b8:0c:cf:d1  
lpfc0t02 DID 6c0b13 WWPN 20:08:00:a0:b8:0b:14:a2 WWNN 20:08:00:a0:b8:0b:14:a1  
lpfc0t03 DID 6c1113 WWPN 20:08:00:a0:b8:0f:55:ff WWNN 20:08:00:a0:b8:0f:55:fe  
lpfc0t04 DID 6c0a13 WWPN 50:05:07:63:00:cf:94:62 WWNN 50:05:07:63:00:c0:94:62
```

```
lpfc0t05 DID 6c1313 WWPN 20:03:00:a0:b8:0f:7f:e4 WWNN 20:02:00:a0:b8:0f:7f:e3
lpfc0t06 DID 6c0713 WWPN 20:04:00:a0:b8:0f:43:63 WWNN 20:04:00:a0:b8:0f:43:62
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2
```

```
GetDriverType: DriverType=00000004
```

```
MULTIPULSE MENU
```

1. Create Route
2. Delete Route
3. Add Paths to Route
4. Delete Paths from Route
5. Route Maintenance
6. Path Maintenance
7. Persistent Bindings

0. Exit MP Menu Application

```
Enter choice => 7
```

```
PERSISTENT BINDINGS MENU
```

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

```
Enter choice => 5
```

```
0. lpfc0 is a LP9002      WWN   : 10:00:00:00:c9:24:0d:a2
1. lpfc1 is a LP9002      WWN   : 10:00:00:00:c9:28:c0:20
2. lpfc2 is a LP9002      WWN   : 10:00:00:00:c9:2b:9e:79
3. lpfc3 is a LP9002      WWN   : 10:00:00:00:c9:24:0c:e3
```

```
Select an adapter (0-3) => 0
```

```
Automapped FCP Nodes:
```

Node	Target	WWPN	WWNN
----	-----	-----	-----
Automapped FCP Node	0	50-05-07-63-00-cc-0e-64	50-05-07-63-00-c0-0e-64
Automapped FCP Node	1	20-02-00-a0-b8-0c-cf-d2	20-02-00-a0-b8-0c-cf-d1
Automapped FCP Node	2	20-08-00-a0-b8-0b-14-a2	20-08-00-a0-b8-0b-14-a1
Automapped FCP Node	3	20-08-00-a0-b8-0f-55-ff	20-08-00-a0-b8-0f-55-fe
Automapped FCP Node	4	50-05-07-63-00-cf-94-62	50-05-07-63-00-c0-94-62
Automapped FCP Node	5	20-03-00-a0-b8-0f-7f-e4	20-02-00-a0-b8-0f-7f-e3
Automapped FCP Node	6	20-04-00-a0-b8-0f-43-63	20-04-00-a0-b8-0f-43-62

```
Bind all auto-mapped targets? (Y/N) => y
```

```
Select Binding Method:
```

1. By Node Name
2. By Port Name
3. By D_ID
0. Cancel

```
Enter Choice => 2
```

```
PERSISTENT BINDINGS MENU
```

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

Enter choice => 5

```

0. lpfc0 is a LP9002   WWN   : 10:00:00:00:c9:24:0d:a2
1. lpfc1 is a LP9002   WWN   : 10:00:00:00:c9:28:c0:20
2. lpfc2 is a LP9002   WWN   : 10:00:00:00:c9:2b:9e:79
3. lpfc3 is a LP9002   WWN   : 10:00:00:00:c9:24:0c:e3

```

Select an adapter (0-3) => 1

Automapped FCP Nodes:

Node	Target	WWPN	WWNN
Automapped FCP Node	0	50-05-07-63-00-c8-94-6d	50-05-07-63-00-c0-94-6d
Automapped FCP Node	1	20-05-00-a0-b8-0f-85-ae	20-04-00-a0-b8-0f-85-ad
Automapped FCP Node	2	20-05-00-a0-b8-0f-7f-12	20-04-00-a0-b8-0f-7f-11

Bind all auto-mapped targets? (Y/N) => y

PERSISTENT BINDINGS MENU

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

Enter choice => 5

```

0. lpfc0 is a LP9002   WWN   : 10:00:00:00:c9:24:0d:a2
1. lpfc1 is a LP9002   WWN   : 10:00:00:00:c9:28:c0:20
2. lpfc2 is a LP9002   WWN   : 10:00:00:00:c9:2b:9e:79
3. lpfc3 is a LP9002   WWN   : 10:00:00:00:c9:24:0c:e3

```

Select an adapter (0-3) => 2

Automapped FCP Nodes:

Node	Target	WWPN	WWNN
Automapped FCP Node	0	10-00-00-00-c9-21-23-f5	50-05-07-63-00-c0-02-ed
Automapped FCP Node	1	20-02-00-a0-b8-0f-7f-e4	20-02-00-a0-b8-0f-7f-e3
Automapped FCP Node	2	20-09-00-a0-b8-0b-14-a2	20-08-00-a0-b8-0b-14-a1
Automapped FCP Node	3	50-05-07-63-00-c6-0e-64	50-05-07-63-00-c0-0e-64
Automapped FCP Node	4	20-03-00-a0-b8-0c-cf-d2	20-02-00-a0-b8-0c-cf-d1
Automapped FCP Node	5	20-05-00-a0-b8-0f-43-63	20-04-00-a0-b8-0f-43-62
Automapped FCP Node	6	20-09-00-a0-b8-0f-55-ff	20-08-00-a0-b8-0f-55-fe
Automapped FCP Node	7	50-05-07-63-00-ce-94-62	50-05-07-63-00-c0-94-62

Bind all auto-mapped targets? (Y/N) => y

PERSISTENT BINDINGS MENU

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

Enter choice => 5

```
0. lpfc0 is a LP9002      WWN : 10:00:00:00:c9:24:0d:a2
1. lpfc1 is a LP9002      WWN : 10:00:00:00:c9:28:c0:20
2. lpfc2 is a LP9002      WWN : 10:00:00:00:c9:2b:9e:79
3. lpfc3 is a LP9002      WWN : 10:00:00:00:c9:24:0c:e3
```

Select an adapter (0-3) => 3

Automapped FCP Nodes:

Node	Target	WWPN	WWNN
Automapped FCP Node	0	20-04-00-a0-b8-0f-85-ae	20-04-00-a0-b8-0f-85-ad
Automapped FCP Node	1	20-04-00-a0-b8-0f-7f-12	20-04-00-a0-b8-0f-7f-11

Bind all auto-mapped targets? (Y/N) => y

PERSISTENT BINDINGS MENU

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 0, WWPN: 50-05-07-63-00-cc-0e-64
2. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
3. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
4. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
5. Adapter: 0, Target: 4, WWPN: 50-05-07-63-00-cf-94-62
6. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
7. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
8. Adapter: 1, Target: 0, WWPN: 50-05-07-63-00-c8-94-6d
9. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
10. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
11. Adapter: 2, Target: 0, WWPN: 10-00-00-00-c9-21-23-f5
12. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
13. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
14. Adapter: 2, Target: 3, WWPN: 50-05-07-63-00-c6-0e-64
15. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
16. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63

17. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
18. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
19. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
20. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

PERSISTENT BINDINGS MENU

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

Enter choice => 6

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings
3. Delete Bindings by Adapter

0. Return to Persistent Bindings Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 0, WWPN: 50-05-07-63-00-cc-0e-64
2. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
3. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
4. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
5. Adapter: 0, Target: 4, WWPN: 50-05-07-63-00-cf-94-62
6. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
7. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
8. Adapter: 1, Target: 0, WWPN: 50-05-07-63-00-c8-94-6d
9. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
10. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
11. Adapter: 2, Target: 0, WWPN: 10-00-00-00-c9-21-23-f5
12. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
13. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
14. Adapter: 2, Target: 3, WWPN: 50-05-07-63-00-c6-0e-64
15. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
16. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63
17. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
18. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
19. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
20. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

Select Binding Number To Delete, or 0 to cancel => 1

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings
3. Delete Bindings by Adapter

0. Return to Persistent Bindings Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
2. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
3. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
4. Adapter: 0, Target: 4, WWPN: 50-05-07-63-00-cf-94-62
5. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
6. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
7. Adapter: 1, Target: 0, WWPN: 50-05-07-63-00-c8-94-6d
8. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
9. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
10. Adapter: 2, Target: 0, WWPN: 10-00-00-00-c9-21-23-f5
11. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
12. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
13. Adapter: 2, Target: 3, WWPN: 50-05-07-63-00-c6-0e-64
14. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
15. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63
16. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
17. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
18. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
19. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

Select Binding Number To Delete, or 0 to cancel => 4

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings
3. Delete Bindings by Adapter
0. Return to Persistent Bindings Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
2. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
3. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
4. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
5. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
6. Adapter: 1, Target: 0, WWPN: 50-05-07-63-00-c8-94-6d
7. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
8. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
9. Adapter: 2, Target: 0, WWPN: 10-00-00-00-c9-21-23-f5
10. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
11. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
12. Adapter: 2, Target: 3, WWPN: 50-05-07-63-00-c6-0e-64
13. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
14. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63
15. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
16. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
17. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
18. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

Select Binding Number To Delete, or 0 to cancel => 6

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings

3. Delete Bindings by Adapter

0. Return to Persistent Bindings Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
2. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
3. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
4. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
5. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
6. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
7. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
8. Adapter: 2, Target: 0, WWPN: 10-00-00-00-c9-21-23-f5
9. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
10. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
11. Adapter: 2, Target: 3, WWPN: 50-05-07-63-00-c6-0e-64
12. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
13. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63
14. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
15. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
16. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
17. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

Select Binding Number To Delete, or 0 to cancel => 8

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings
3. Delete Bindings by Adapter

0. Return to Persistent Bindings Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
2. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
3. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
4. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
5. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
6. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
7. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
8. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
9. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
10. Adapter: 2, Target: 3, WWPN: 50-05-07-63-00-c6-0e-64
11. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
12. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63
13. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
14. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
15. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
16. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

Select Binding Number To Delete, or 0 to cancel => 10

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings

3. Delete Bindings by Adapter

0. Return to Persistent Bindings Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
2. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
3. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
4. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
5. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
6. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
7. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
8. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
9. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
10. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
11. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63
12. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
13. Adapter: 2, Target: 7, WWPN: 50-05-07-63-00-ce-94-62
14. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
15. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12

Select Binding Number To Delete, or 0 to cancel => 13

DELETE BINDINGS MENU

1. Delete Specific Binding
2. Delete All Bindings
3. Delete Bindings by Adapter
0. Return to Persistent Bindings Menu

Enter choice => 0

PERSISTENT BINDINGS MENU

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)
0. Return to Main Menu

Enter choice => 1

Persistent Bindings:

1. Adapter: 0, Target: 1, WWPN: 20-02-00-a0-b8-0c-cf-d2
2. Adapter: 0, Target: 2, WWPN: 20-08-00-a0-b8-0b-14-a2
3. Adapter: 0, Target: 3, WWPN: 20-08-00-a0-b8-0f-55-ff
4. Adapter: 0, Target: 5, WWPN: 20-03-00-a0-b8-0f-7f-e4
5. Adapter: 0, Target: 6, WWPN: 20-04-00-a0-b8-0f-43-63
6. Adapter: 1, Target: 1, WWPN: 20-05-00-a0-b8-0f-85-ae
7. Adapter: 1, Target: 2, WWPN: 20-05-00-a0-b8-0f-7f-12
8. Adapter: 2, Target: 1, WWPN: 20-02-00-a0-b8-0f-7f-e4
9. Adapter: 2, Target: 2, WWPN: 20-09-00-a0-b8-0b-14-a2
10. Adapter: 2, Target: 4, WWPN: 20-03-00-a0-b8-0c-cf-d2
11. Adapter: 2, Target: 5, WWPN: 20-05-00-a0-b8-0f-43-63

```
12. Adapter: 2, Target: 6, WWPN: 20-09-00-a0-b8-0f-55-ff
13. Adapter: 3, Target: 0, WWPN: 20-04-00-a0-b8-0f-85-ae
14. Adapter: 3, Target: 1, WWPN: 20-04-00-a0-b8-0f-7f-12
```

PERSISTENT BINDINGS MENU

1. Display Current Bindings
2. Display All Nodes
3. Duplicate Adapter Bindings
4. Bind Target Manually
5. Bind Automapped Targets
6. Delete Binding(s)

0. Return to Main Menu

```
Enter choice => 0
```

MULTIPULSE MENU

1. Create Route
2. Delete Route
3. Add Paths to Route
4. Delete Paths from Route
5. Route Maintenance
6. Path Maintenance
7. Persistent Bindings

0. Exit MP Menu Application

```
Enter choice => 0
```

4. Single path configuration for FAStT Multipulse devices

Edit `/etc/lpfc.conf` to change `lpfc_automap=0` and `lpfc_scandown=0`; to build a single path driver with only FAStT Multipulse devices in Non-Multipulse configuration.

```
# vi /etc/lpfc.conf // change lpfc_automap=0 and lpfc_scandown=0
# make build >/dev/null 2>/dev/null
# make install
Installing under /lib/modules/2.4.21-111-pseries64/kernel/drivers
```

Load the new driver for a single path configuration with FAStT Multipulse devices.

```
# rmmod lpfcdd
# modprobe lpfcdd
Warning: loading /lib/modules/2.4.21-111-
pseries64/kernel/drivers/scsi/lpfcdd.o will taint the kernel: no license
See http://www.tux.org/lkml/#export-tainted for information about
tainted modules
Module lpfcdd loaded, with warnings
```

5. Create entries for FASSt Multipulse devices in multi-path configurations

Here is an example to create a Multipulse entry for a multi-path device with autofailback.

```
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2
GetDriverType: DriverType=00000004
MULTIPULSE MENU

1. Create Route
2. Delete Route
3. Add Paths to Route
4. Delete Paths from Route
5. Route Maintenance
6. Path Maintenance
7. Persistent Bindings

0. Exit MP Menu Application

Enter choice => 1
 1. lpfc0t1d0 600A0B80000CCFD1000000013EFB0DB2
 2. lpfc0t1d1 600A0B80000CD094000000003EFB0E5F
 3. lpfc0t1d2 600A0B80000CCFD1000000233FA92D44
 4. lpfc0t2d0 600A0B8000075B440000024E3EA6CBD4
 5. lpfc0t3d1 600A0B80000F55FE000003D03EDF2DF6
 6. lpfc0t5d0 600A0B80000F7FE3000000223FA7E167
 7. lpfc0t5d1 600A0B80000F8046000000323FA7E16A
 8. lpfc0t6d0 600A0B80000F43620000001B3EA3FF16
 9. lpfc0t6d1 600A0B80000F4356000000273EA3FE17
10. lpfc0t6d2 600A0B80000F4362000000483EA41D4A
11. lpfc0t6d3 600A0B80000F4362000000493EA41D4E
12. lpfc0t6d4 600A0B80000F43620000004A3EA41D52
13. lpfc0t6d5 600A0B80000F43620000004B3EA41D56
14. lpfc0t6d6 600A0B80000F4362000000193EA3FEB6
15. lpfc0t6d7 600A0B80000F4356000000263EA3FDB7
16. lpfc1t1d0 600A0B80000F876F000005F43F72B2C8

0. Return to MultiPulse Main Menu

Enter 's' or 'S' to continue scrolling display for target choices.

Select a target => 1
 1. lpfc0t1d0
 2. lpfc2t4d0

Select paths (maximum 4 path; 0=done) =>1 2 0
AutoFailBack ? (Y/N) => y
InquiryHeartbeat ? (Y/N) => n
FailoverAbortTaskSet ? (Y/N) => n
FailoverLunReset ? (Y/N) => n
FailoverTgtReset ? (Y/N) => n

Control type :
 1. Simple Failover
 2. Traffic Ratio
 3. Lun Load Balancing
 4. Target Load Balancing
```

5. HBA Load Balancing

Select a type (1 - 5) => 1

1. lpfc0t1d0
2. lpfc2t4d0

Select primary path => 1

Are you sure you want to create this route ? (Y/N) => y

MULTIPULSE MENU

1. Create Route
2. Delete Route
3. Add Paths to Route
4. Delete Paths from Route
5. Route Maintenance
6. Path Maintenance
7. Persistent Bindings

0. Exit MP Menu Application

Enter choice => 0

Here is an example to create Multipulse entries for multi-path devices with autofailback using template (“addroute” with 10 Multipulse entries).

```
# cat addroute
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
1 1 1 2 0 y n n n n 1 1 y
0
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2 < addroute
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2 < addroute
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2 < addroute
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2 < addroute
# /usr/sbin/hbanyware/hbacmd mpMenu 10:00:00:00:c9:24:0d:a2 < addroute
```

Here is the output of /etc/lpfc.conf after the configuration change.

```
# cat /etc/lpfc.conf
/*
 * COPYRIGHT 2003, EMULEX CORPORATION
 * 3535 Harbor Boulevard, Costa Mesa, CA 92626
 *
 * This computer program and related documentation ("Program") is provided
 * under a non-exclusive, worldwide, royalty-free, perpetual, transferable
 * license to use, copy, compile, modify or have modified, display,
 * distribute or have distributed, with the exception of limitations
 * associated with the subsidiary program ("Subsidiary") contained therein.
 * The Subsidiary will be provided in object code only and may not be
 * modified, displayed, decompiled, or reverse engineered. The Subsidiary
 * is confidential and a trade secret of EMULEX CORPORATION. Any distribution
 * of the Program without this copyright notice is a violation of the
 * copyright laws and may subject you to criminal prosecution. THE PROGRAM
 * IS LICENSED AS IS, WITHOUT ANY WARRANTY, INCLUDING THE IMPLIED WARRANTIES
 * OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, TITLE, AND AGAINST
 * INFRINGEMENT.
 *
 * $Id: lpfc.conf 1.24 2002/06/03 16:08:24 mks Exp $
 *
 * LINUX LightPulse lpfc (SCSI) / lpfn (IP) driver: global initialized data.
 */

/*
# Setup FCP persistent bindings,
# fcp-bind-WWPN binds a specific WorldWide PortName to a target id,
# fcp-bind-WWNN binds a specific WorldWide NodeName to a target id,
# fcp-bind-DID binds a specific DID to a target id.
# Only one binding method can be used. "lpfc_automap" needs to
# be changed to 0 and scan-down should NOT be set to 2
# when one of these binding methods is used.
# WWNN, WWPN and DID are hexadecimal values.
# WWNN must be 16 digit BCD with leading 0s.
# WWPN must be 16 digit BCD with leading 0s.
# DID must be 6 digit BCD with leading 0s.
# The SCSI ID to bind to consists of two parts, the lpfc interface
# to bind to, and the target number for that interface.
# Thus lpfc0t2 specifies target 2 on interface lpfc0.
#
# Here are some examples:
#
#                               WWNN           SCSI ID
# char *lpfc_fcp_bind_WWNN[]={ "22000020370b8275:lpfc0t1",
#                               "22000020370b8998:lpfc0t2"};
#
#                               WWPN           SCSI ID
# char *lpfc_fcp_bind_WWPN[]={ "22000020370b8275:lpfc0t1",
#                               "22000020370b8998:lpfc0t2"};
#
#                               DID    SCSI ID
# char *lpfc_fcp_bind_DID[]={ "0000dc:lpfc0t1",
#                               "0000e0:lpfc0t2"};
#
 */
int    lpfc_bind_entries=14;

char    *lpfc_fcp_bind_WWNN[MAX_FC_BINDINGS];
```



```

char    *lpfc_fcp_bind_WWPN[MAX_FC_BINDINGS]={
    "200200a0b80ccfd2:lpfc0t1",
    "200800a0b80b14a2:lpfc0t2",
    "200800a0b80f55ff:lpfc0t3",
    "200300a0b80f7fe4:lpfc0t5",
    "200400a0b80f4363:lpfc0t6",
    "200500a0b80f85ae:lpfc1t1",
    "200500a0b80f7f12:lpfc1t2",
    "200200a0b80f7fe4:lpfc2t1",
    "200900a0b80b14a2:lpfc2t2",
    "200300a0b80ccfd2:lpfc2t4",
    "200500a0b80f4363:lpfc2t5",
    "200900a0b80f55ff:lpfc2t6",
    "200400a0b80f85ae:lpfc3t0",
    "200400a0b80f7f12:lpfc3t1"};

char    *lpfc_fcp_bind_DID[MAX_FC_BINDINGS];

/*
# If automap is set, SCSI IDs for all FCP nodes without
# persistent bindings will be automatically generated.
# If new FCP devices are added to the network when the system is down,
# there is no guarantee that these SCSI IDs will remain the same
# when the system is booted again.
# If one of the above fcp binding methods is specified, then automap
# devices will use the same mapping method to preserve
# SCSI IDs between link down and link up.
# If no bindings are specified above, a value of 1 will force WWNN
# binding, 2 for WWPN binding, and 3 for DID binding.
# If automap is 0, only devices with persistent bindings will be
# recognized by the system.
*/
int    lpfc_automap=0;

/*
# Some disk devices have a "select ID" or "select Target" capability.
# From a protocol standpoint "select ID" usually means select the
# Fibre channel "ALPA". In the FC-AL Profile there is an "informative
# annex" which contains a table that maps a "select ID" (a number
# between 0 and 7F) to an ALPA. By default, for compatibility with
# older drivers, the lpfc driver scans its ALPA map from low ALPA
# to high ALPA.
#
# Turning on the scan-down variable (on = 1, off = 0) will
# cause the lpfc driver to use an inverted ALPA map, effectively
# scanning ALPAs from high to low as specified in the FC-AL annex.
# A value of 2, will also cause target assignment in a private loop
# environment to be based on the ALPA. Persistent bindings should NOT be
# used if scan-down is 2.
#
# (Note: This "select ID" functionality is a LOOP ONLY characteristic
# and will not work across a fabric.)
*/
int    lpfc_scandown=0;

#ifdef MPULSE
/*
# MultiPulse configuration

```

```

#
# multipulse-fcp="lpfcXtYdZ:tr1: ... :lpfcXtYdZ:trN:route_flags";
# Where X, Y, and Z are the devices lpfc DDI interface, target, and LUN
# numbers.
# tr1 thru trN, when N is up to 4, are the paths component of the traffic
# ratio
# route_flags MUST be a 4 digit hex number for the following flags:
# Load balancing flags (1 means cold standby)
# MPL_PC_TYPE_FAILOVER          0x0001  cold standby
# MPL_PC_TYPE_TRAFFIC_RATIO     0x0002  paths balanced by traffic ratio
# MPL_PC_TYPE_DYNAMIC_LUNQ      0x0003  Dynamic balancing algorithm used
# MPL_PC_TYPE_DYNAMIC_TGTQ      0x0004  Dynamic balancing algorithm used
# MPL_PC_TYPE_DYNAMIC_HBAQ      0x0005  Dynamic balancing algorithm used
#
# General flags
# MPL_PC_CFG_AUTO_FAILBACK      0x0010  failback to primary path is auto
# MPL_PC_CFG_INQUIRY_HEARTBEAT  0x0040  Validate path inquiry heartbeat
# MPL_PC_CFG_FAIL_ABORT_TASK_SET 0x0080  Send ABORT_TASK_SET on failover
# MPL_PC_CFG_FAIL_LUN_RESET      0x0100  Send LUN_RESET on failover
# MPL_PC_CFG_FAIL_TARGET_RESET  0x0200  Send TARGET_RESET on failover
#
# The following example uses cold standby as path control and automatic
# failback
#
# Here is a sample configuration to setup lpfc0 target 0 lun 1 failover
# to the MultiPulse device, lpfc1 target 0 lun 1 and lpfc0 target 0 lun 2
# failover to the MultiPulse device, lpfc1 target 1 lun 4.
# char *lpfc_multipulse_fcp[]={"lpfc0t0d1:0:lpfc1t0d1:0:0010",
#                               "lpfc0t0d2:0:lpfc1t1d4:0:0010"};
#
# NOTE: When using multipulse it is highly recommended that linkdown_tmo
# and nodev_tmo be reduced to 1 or 2 seconds. These two timeouts will directly
# effect failover time.
#
# BEGIN: MultiPulse managed entries
*/
int    lpfc_multipulse_entries=47;

char   *lpfc_multipulse_fcp[47]={
        "lpfc0t1d0:0:lpfc2t4d0:0:0011",
        "lpfc0t1d1:0:lpfc2t4d1:0:0011",
        "lpfc0t1d2:0:lpfc2t4d2:0:0011",
        "lpfc0t2d0:0:lpfc2t2d0:0:0011",
        "lpfc0t3d1:0:lpfc2t6d1:0:0011",
        "lpfc0t5d0:0:lpfc2t1d0:0:0011",
        "lpfc0t5d1:0:lpfc2t1d1:0:0011",
        "lpfc0t6d0:0:lpfc2t5d0:0:0011",
        "lpfc0t6d1:0:lpfc2t5d1:0:0011",
        "lpfc0t6d2:0:lpfc2t5d2:0:0011",
        "lpfc0t6d3:0:lpfc2t5d3:0:0011",
        "lpfc0t6d4:0:lpfc2t5d4:0:0011",
        "lpfc0t6d5:0:lpfc2t5d5:0:0011",
        "lpfc0t6d6:0:lpfc2t5d6:0:0011",
        "lpfc0t6d7:0:lpfc2t5d7:0:0011",
        "lpfc1t1d0:0:lpfc3t0d0:0:0011",
        "lpfc1t1d1:0:lpfc3t0d1:0:0011",
        "lpfc1t1d2:0:lpfc3t0d2:0:0011",
        "lpfc1t1d3:0:lpfc3t0d3:0:0011",
        "lpfc1t1d4:0:lpfc3t0d4:0:0011",

```

```
    "lpfc1t1d5:0:lpfc3t0d5:0:0011",
    "lpfc1t1d6:0:lpfc3t0d6:0:0011",
    "lpfc1t1d7:0:lpfc3t0d7:0:0011",
    "lpfc1t1d8:0:lpfc3t0d8:0:0011",
    "lpfc1t1d9:0:lpfc3t0d9:0:0011",
    "lpfc1t1d10:0:lpfc3t0d10:0:0011",
    "lpfc1t1d11:0:lpfc3t0d11:0:0011",
    "lpfc1t1d12:0:lpfc3t0d12:0:0011",
    "lpfc1t1d13:0:lpfc3t0d13:0:0011",
    "lpfc1t1d14:0:lpfc3t0d14:0:0011",
    "lpfc1t1d15:0:lpfc3t0d15:0:0011",
    "lpfc1t2d0:0:lpfc3t1d0:0:0011",
    "lpfc1t2d1:0:lpfc3t1d1:0:0011",
    "lpfc1t2d2:0:lpfc3t1d2:0:0011",
    "lpfc1t2d3:0:lpfc3t1d3:0:0011",
    "lpfc1t2d4:0:lpfc3t1d4:0:0011",
    "lpfc1t2d5:0:lpfc3t1d5:0:0011",
    "lpfc1t2d6:0:lpfc3t1d6:0:0011",
    "lpfc1t2d7:0:lpfc3t1d7:0:0011",
    "lpfc1t2d8:0:lpfc3t1d8:0:0011",
    "lpfc1t2d9:0:lpfc3t1d9:0:0011",
    "lpfc1t2d10:0:lpfc3t1d10:0:0011",
    "lpfc1t2d11:0:lpfc3t1d11:0:0011",
    "lpfc1t2d12:0:lpfc3t1d12:0:0011",
    "lpfc1t2d13:0:lpfc3t1d13:0:0011",
    "lpfc1t2d14:0:lpfc3t1d14:0:0011",
    "lpfc1t2d15:0:lpfc3t1d15:0:0011"
};
#endif
```

6. Build a Multipulse driver with FAStT multi-path devices

```
# make build >/dev/null 2>/dev/null
# make install
Installing under /lib/modules/2.4.21-111-pseries64/kernel/drivers
```

7. Final Step – load the new driver for multi-path configuration

```
# modprobe lpfcdd 1>/dev/null 2>/dev/null
# scsidev
# cat /proc/scsi/scsi
Attached devices:
Host: scsi0 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: CDRM00203    !K Rev: 1_05
  Type:   CD-ROM      ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 04 Lun: 00
  Vendor: IBM      Model: DDYS-T09170N   Rev: S9RA
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 1742           Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 01
  Vendor: IBM      Model: 1742           Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: IBM      Model: 1742           Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 02 Lun: 00
  Vendor: IBM      Model: 3542           Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 03 Lun: 00
  Vendor: IBM      Model: 1722-600   Rev: 0520
  Type:   Processor    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 03 Lun: 01
  Vendor: IBM      Model: 1722-600   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 05 Lun: 00
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 05 Lun: 01
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 00
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 01
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 02
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 03
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 04
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 05
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 06
  Vendor: IBM      Model: 1742-900   Rev: 0520
  Type:   Direct-Access ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 06 Lun: 07
```

```

Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 00
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 01
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 02
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 03
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 04
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 05
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 06
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 07
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 08
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 09
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 10
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 11
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 12
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 13
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 14
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 15
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 00
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 01
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 02
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03

```

```

Host: scsi3 Channel: 00 Id: 02 Lun: 03
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 04
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 05
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 06
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 07
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 08
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 09
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 10
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 11
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 12
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 13
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 14
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 02 Lun: 15
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 01 Lun: 00
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 01 Lun: 01
Vendor: IBM      Model: 1742-900      Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 02 Lun: 00
Vendor: IBM      Model: 3542          Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 04 Lun: 00
Vendor: IBM      Model: 1742          Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 04 Lun: 01
Vendor: IBM      Model: 1742          Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 04 Lun: 02
Vendor: IBM      Model: 1742          Rev: 0520
Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 00
Vendor: IBM      Model: 1742-900      Rev: 0520

```

```

Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 01
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 02
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 03
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 04
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 05
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 06
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 05 Lun: 07
Vendor: IBM Model: 1742-900 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 06 Lun: 00
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Processor ANSI SCSI revision: 03
Host: scsi4 Channel: 00 Id: 06 Lun: 01
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 00
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 01
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 02
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 03
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 04
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 05
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 06
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 07
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 08
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 09
Vendor: IBM Model: 1722-600 Rev: 0520
Type: Direct-Access ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 10

```

```

Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 11
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 12
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 13
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 14
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 00 Lun: 15
Vendor: IBM      Model: 1722-600      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 00
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 01
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 02
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 03
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 04
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 05
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 06
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 07
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 08
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 09
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 10
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 11
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 12
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 13
Vendor: IBM      Model: 1742-900      Rev: 0520
Type:   Direct-Access                ANSI SCSI revision: 03

```



```

Host: scsi5 Channel: 00 Id: 01 Lun: 14
  Vendor: IBM      Model: 1742-900      Rev: 0520
  Type:   Direct-Access      ANSI SCSI revision: 03
Host: scsi5 Channel: 00 Id: 01 Lun: 15
  Vendor: IBM      Model: 1742-900      Rev: 0520
  Type:   Direct-Access      ANSI SCSI revision: 03
# cat /proc/scsi/lpfc/2
Emulex LightPulse LPFC Driver Version: 4.30o_M
HBA: Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
SerialNum: 0000c9240da2
Firmware Version: 3.91A1
Hdw: 1002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:24:0d:a2  Nodename: 20:00:00:00:c9:24:0d:a2

Link Up - Ready:
  PortID 0x6c1213
  Fabric
  Current speed 2G

lpfc0t01 DID 6c0c13 WWPN 20:02:00:a0:b8:0c:cf:d2 WWNN 20:02:00:a0:b8:0c:cf:d1
lpfc0t02 DID 6c0b13 WWPN 20:08:00:a0:b8:0b:14:a2 WWNN 20:08:00:a0:b8:0b:14:a1
lpfc0t03 DID 6c1113 WWPN 20:08:00:a0:b8:0f:55:ff WWNN 20:08:00:a0:b8:0f:55:fe
lpfc0t05 DID 6c1313 WWPN 20:03:00:a0:b8:0f:7f:e4 WWNN 20:02:00:a0:b8:0f:7f:e3
lpfc0t06 DID 6c0713 WWPN 20:04:00:a0:b8:0f:43:63 WWNN 20:04:00:a0:b8:0f:43:62
# cat /proc/scsi/lpfc/3
Emulex LightPulse LPFC Driver Version: 4.30o_M
HBA: Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
SerialNum: 0000c928c020
Firmware Version: 3.91A1
Hdw: 2002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:28:c0:20  Nodename: 20:00:00:00:c9:28:c0:20

Link Up - Ready:
  PortID 0x422100
  Fabric
  Current speed 2G

lpfc1t01 DID 422700 WWPN 20:05:00:a0:b8:0f:85:ae WWNN 20:04:00:a0:b8:0f:85:ad
lpfc1t02 DID 422200 WWPN 20:05:00:a0:b8:0f:7f:12 WWNN 20:04:00:a0:b8:0f:7f:11
# cat /proc/scsi/lpfc/4
Emulex LightPulse LPFC Driver Version: 4.30o_M
HBA: Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
SerialNum: 1421R01064
Firmware Version: 3.91A1
Hdw: 1002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:2b:9e:79  Nodename: 20:00:00:00:c9:2b:9e:79

Link Up - Ready:
  PortID 0xa0a00
  Fabric
  Current speed 2G

lpfc2t01 DID 0a0700 WWPN 20:02:00:a0:b8:0f:7f:e4 WWNN 20:02:00:a0:b8:0f:7f:e3
lpfc2t02 DID 0a0600 WWPN 20:09:00:a0:b8:0b:14:a2 WWNN 20:08:00:a0:b8:0b:14:a1
lpfc2t04 DID 0a0300 WWPN 20:03:00:a0:b8:0c:cf:d2 WWNN 20:02:00:a0:b8:0c:cf:d1
lpfc2t05 DID 0a0200 WWPN 20:05:00:a0:b8:0f:43:63 WWNN 20:04:00:a0:b8:0f:43:62

```

```
lpfc2t06 DID 0a0100 WWPN 20:09:00:a0:b8:0f:55:ff WWNN 20:08:00:a0:b8:0f:55:fe
# cat /proc/scsi/lpfc/5
Emulex LightPulse LPFC Driver Version: 4.30o_M
HBA: Emulex LightPulse LP9002 2 Gigabit PCI Fibre Channel Adapter
SerialNum: 0000c9240ce3
Firmware Version: 3.91A1
Hdw: 1002606d
VendorId: 0xf90010df
Portname: 10:00:00:00:c9:24:0c:e3   Nodename: 20:00:00:00:c9:24:0c:e3

Link Up - Ready:
  PortID 0x50301
  Fabric
  Current speed 2G

lpfc3t00 DID 050200 WWPN 20:04:00:a0:b8:0f:85:ae WWNN 20:04:00:a0:b8:0f:85:ad
lpfc3t01 DID 050100 WWPN 20:04:00:a0:b8:0f:7f:12 WWNN 20:04:00:a0:b8:0f:7f:11
```



© IBM Corporation 2004

IBM Corporation
Marketing Communications
Systems Group
Route 100
Somers, New York 10589

Produced in the United States of America
January 2004
All Rights Reserved

This document was developed for products and/or services offered in the United States. IBM may not offer the products, features, or services discussed in this document in other countries. The information may be subject to change without notice. Consult your local IBM business contact for information on the products, features and services available in your area.

All statements regarding IBM's future directions and intent are subject to change or withdrawal without notice and represent goals and objectives only.

IBM, the IBM logo, the e-business logo, @server, AIX 5L, iSeries, PowerPC, PowerPC Architecture, pSeries, and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both. A full list of U.S. trademarks owned by IBM may be found at <http://www.ibm.com/legal/copytrade.shtml>.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Windows is a registered trademark of the Microsoft Corporation.

Intel and Itanium are registered trademarks of Intel Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.

IBM hardware products are manufactured from new parts, or new and used parts. Regardless, our warranty terms apply.

This equipment is subject to FCC rules. It will comply with the appropriate FCC rules before final delivery to the buyer.

Information concerning non-IBM products was obtained from the suppliers of these products or other public sources. Questions on the capabilities of the non-IBM products should be addressed with the suppliers.

All performance information was determined in a controlled environment. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

The IBM home page on the Internet can be found at <http://www.ibm.com>.

The pSeries home page on the Internet can be found at <http://www.ibm.com/servers/eserver/pSeries>.