



EMULEX[®]

We network storage

Emulex Driver for Solaris

Version 6.02h

User Manual

Copyright© 2005 Emulex Corporation. All rights reserved worldwide. No part of this document may be reproduced by any means nor translated to any electronic medium without the written consent of Emulex Corporation.

Information furnished by Emulex Corporation is believed to be accurate and reliable. However, no responsibility is assumed by Emulex Corporation for its use; or for any infringements of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under any patent or patent rights of Emulex Corporation.

Emulex and LightPulse are registered trademarks, and AutoPilot Installer, AutoPilot Manager, BlockGuard, FibreSpy, HBAnyware, InSpeed, MultiPulse and SBOD are trademarks, of Emulex Corporation. All other brand or product names referenced herein are trademarks or registered trademarks of their respective companies or organizations.

Emulex provides this manual "as is" without any warranty of any kind, either expressed or implied, including but not limited to the implied warranties of merchantability or fitness for a particular purpose. Emulex Corporation may make improvements and changes to the product described in this manual at any time and without any notice. Emulex Corporation assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties that may result. Periodic changes are made to information contained herein; although these changes will be incorporated into new editions of this manual, Emulex Corporation disclaims any undertaking to give notice of such changes.

Installation 1

| | |
|--|----|
| Driver Information | 1 |
| Supported Features..... | 1 |
| New Features in this Release..... | 2 |
| Prerequisites | 2 |
| Compatibility..... | 2 |
| Things to Know Before You Download | 3 |
| Known Issues | 4 |
| Installing the Emulex Driver for Solaris Version 6.02h | 6 |
| Introduction | 6 |
| Installing the Driver for the First Time..... | 6 |
| Installing the Driver as a Loadable Kernel Module | 8 |
| Updating Existing Drivers | 9 |
| Updating a PCI Driver (from Version 6.X)..... | 9 |
| Updating a PCI Driver (from Version 5.X)..... | 10 |
| update_lpfcb | 11 |
| Updating a PCI Driver (from Version 4.X)..... | 11 |
| Updating an SBus Driver (from Version 4.X) | 12 |
| Updating PCI and SBus Drivers (from Version 4.X) in One System | 13 |
| Directory Structure | 15 |
| Installing the Utilities..... | 15 |
| Installing HBAnyware and lputil | 15 |
| Prerequisites | 15 |
| Procedure | 15 |
| Installing the HBAnyware Security Configurator | 16 |
| Prerequisites | 16 |
| Procedure | 16 |
| Removing an Existing Driver | 16 |
| Uninstalling the Utilities..... | 17 |
| Uninstalling the HBAnyware Security Configurator..... | 17 |
| Uninstalling HBAnyware and lputil..... | 17 |
| Loading Manpages | 17 |
| SUN SAN Foundation Software Considerations | 17 |

Configuration 19

| | |
|--|----|
| Introduction..... | 19 |
| Starting HBAnyware for Solaris | 20 |
| Starting HBAnyware in Local Element Manager Mode | 20 |
| HBAnyware Window Element Definitions | 21 |
| The Menu Bar | 21 |
| The Toolbar | 21 |
| The Toolbar Icon Definitions..... | 22 |
| Sort and Display Icons | 22 |
| The Discovery Tree | 23 |
| Property Tabs..... | 24 |
| Status Bar | 24 |
| Using the HBAnyware Command-Line Interface | 24 |
| Starting lputil..... | 27 |

| | |
|---|----|
| No-Boot Features | 27 |
| Discovering New Targets and LUNS | 27 |
| Removing Ghost LUNS | 29 |
| Adding New Adapters, Changing Bindings and Changing Driver Parameters | 29 |
| Turning on Logging Capability | 30 |
| No-Boot Firmware Updates | 30 |
| Loading or Unloading the Driver Without Booting..... | 31 |
| HBA Tasks | 32 |
| Discovering HBAs | 32 |
| Discovering HBAs Using HBAnyware | 32 |
| Discovering HBAs Using lputil..... | 33 |
| Sorting HBA Information..... | 34 |
| Sorting HBAs Using HBAnyware..... | 34 |
| Sorting Local HBAs Only Using HBAnyware | 35 |
| Sorting Local HBAs Using lputil..... | 35 |
| Viewing HBA Information Using HBAnyware | 35 |
| Viewing Discovered Elements | 35 |
| Viewing Host Attributes | 36 |
| Viewing Target Attributes | 37 |
| Viewing LUN Attributes | 38 |
| Viewing Fabric Attributes | 40 |
| Viewing General HBA Attributes | 41 |
| Viewing Detailed HBA Attributes | 42 |
| Viewing Port Information | 44 |
| Viewing Statistics | 45 |
| Viewing Firmware Information..... | 47 |
| Viewing Target Mapping | 48 |
| Viewing Driver Parameters..... | 49 |
| Viewing HBA Information using lputil..... | 52 |
| Resetting Adapters | 53 |
| Resetting the HBA Using HBAnyware | 53 |
| Resetting the HBA Using lputil | 53 |
| Updating Firmware | 54 |
| Updating Firmware Using HBAnyware | 54 |
| Prerequisites | 54 |
| Procedure | 54 |
| Loading Firmware (Batch Mode) Using HBAnyware..... | 55 |
| Prerequisites | 55 |
| Procedure | 55 |
| Updating Firmware Using lputil..... | 57 |
| Prerequisites | 57 |
| Procedure | 57 |
| Updating FC Boot | 57 |
| Updating FC Boot Code Using HBAnyware..... | 57 |
| Prerequisites | 57 |
| Procedure | 58 |
| Updating FC Boot Code Using lputil..... | 58 |
| Prerequisites | 58 |
| Procedure | 58 |
| Enabling/Disabling BootBIOS Using lputil | 58 |
| Prerequisites | 58 |

| | |
|---|----|
| Procedure | 59 |
| Configuring the Driver..... | 59 |
| Configuring Driver Parameters Using HBAnyware | 59 |
| Change a Parameter's Value..... | 59 |
| Restoring All Parameters to Their Earlier Values | 60 |
| Resetting All Default Values | 60 |
| Configuring Driver Parameters Using lpfc.conf | 61 |
| Using lpfc.conf for Solaris 9 and 10..... | 61 |
| Updating Parameters for Solaris 2.6, 7 and 8 | 61 |
| Understanding Device Numbering..... | 62 |
| The Configuration File (lpfc.conf) | 62 |
| Probing for FCP Disk Targets -- Configuring sd.conf | 73 |
| Probing for FCP Tape Targets -- Configuring st.conf..... | 74 |
| Viewing Target Mapping and Set Up Persistent Binding | 75 |
| Viewing Target Mapping Using HBAnyware | 75 |
| Setting Up Persistent Binding Using HBAnyware | 75 |
| Adding New Targets Using sd.conf (Solaris 8)..... | 78 |
| Changing Parameters or Bindings..... | 78 |
| Setting Up Persistent Binding using lputil | 78 |
| Setting Up Target/LUN Blocking Using sd.conf..... | 79 |
| Enabling IP Networking | 81 |
| Overview | 81 |
| Enabling the Networking Driver Parameter..... | 81 |
| Creating Hostname.lpfc# Files | 81 |
| Required Hostname File | 81 |
| Additional Hostname Files | 81 |
| Verifying the Network Connection | 82 |
| Enabling IP Networking for Multiple Adapters..... | 82 |
| Creating Additional Hostname.lpfc# Files | 82 |
| Configuring the System-Wide File (/etc/system) | 83 |
| Driver Installed in a Non-Standard Directory (moddir) | 83 |
| IP Networking Enabled (forceload)..... | 83 |
| HBAnyware Security..... | 84 |
| Introduction | 84 |
| Starting the HBAnyware Security Configurator | 84 |
| Prerequisites | 84 |
| Procedure | 85 |
| Running the Configurator for the First Time/Creating the ACG..... | 85 |
| Designating a Master Security Client..... | 86 |
| Access Control Groups..... | 86 |
| Introduction | 86 |
| Access Control Group Tab on the MSC..... | 86 |
| Access Control Group Tab on a Non-MSC | 87 |
| ACG Icons..... | 88 |
| Creating the ACG..... | 88 |
| Adding a Server to the ACG..... | 89 |
| Deleting a Server from the ACG..... | 90 |
| Removing Security from all Servers in the ACG..... | 91 |
| Generating New Security Keys | 93 |
| Restoring the ACG to Its Last Saved Configuration | 94 |
| Accessing a Switch | 94 |

| | |
|---|-----|
| Accessing Sub-Groups | 95 |
| Introduction | 95 |
| ASG Icons | 96 |
| Creating an ASG | 96 |
| Reserved Indices - Examples | 98 |
| Adding a Server to an ASG | 98 |
| Deleting an ASG | 99 |
| Restoring an ASG to Its Last Saved Configuration | 99 |
| Editing an ASG | 99 |
| About Offline ASGs | 101 |
| Backup Masters | 101 |
| Introduction | 101 |
| Backup Master Eligible Systems | 102 |
| Backup Master Tab and Controls | 103 |
| Creating a Backup Master | 103 |
| Reassigning a Backup Master as the New MSC from the Old MSC | 104 |
| Reassigning a Backup Master as the New MSC from the Backup Master | 104 |

Troubleshooting..... 106

| | |
|--|-----|
| Introduction..... | 106 |
| Unusual Situations and Their Resolutions | 106 |
| General Situations | 106 |
| Security Configurator Situations - Access Control Groups (ACG) | 108 |
| Security Configuration Situations - Access Sub-Groups (ASG) | 109 |
| HBAware Security Configurator Situations - Backup Masters | 110 |
| Error Message Situations | 111 |
| Master Security Client Situations..... | 112 |
| Non-Hierarchical and Hierarchical ASG | 113 |
| Ipfc Log Messages | 114 |
| Introduction | 114 |
| Severity Codes..... | 115 |
| Message Group Masks | 115 |
| Message Log Example..... | 116 |
| ELS Events (0100 - 0199) | 116 |
| Link Discovery Events (0200 - 0299)..... | 120 |
| Mailbox Events (0300 - 0399)..... | 127 |
| Initialization Events (0400 - 0499) | 132 |
| IP Traffic History (0600 - 0699)..... | 139 |
| FCP Traffic History (0700 - 0799)..... | 141 |
| Node Table Events (0900 - 0999)..... | 146 |
| Miscellaneous Events (1200 - 1299) | 149 |
| Link Events (1300 - 1399) | 150 |
| Log Messages - IOCTL Events (1600 - 1699) | 151 |

Installation

Driver Information

Supported Features

- Supports 256 LUNs (0-255)
- Supports dynamically adding LUNs and targets
- Simplified driver installation
- Topology support: FC-AL, point-to-point, fabric with auto-topology negotiation
- Support for 1, 2 and 4 Gb with auto rate negotiation
- Protocols: SCSI/FC, IP/FC
- Persistent bindings by WWNN, WWPN or D_ID (different methods can be set on an adapter port basis)
- Support for up to thirty-two host bus adapter ports
- Monitoring, parameter configuration and binding settings using Emulex's HBAnyware™ Java-based graphical user interface utility
- Parameter configuration using Emulex's LightPulse® lputil command-line interface utility
- Support for Common HBA API
- lpfc.conf migration script. This script enables you to convert 5.0x lpfc.conf syntax to 6.x lpfc.conf syntax. The script, update_lpfc, is bundled with the driver kit and is installed in the /usr/sbin/lpfc directory.
- Dynamic target/LUN discovery (HBA automatically discovers new targets - Administrator assigns SCSI ID using HBAnyware).
- Dynamic parameter setting (see parameter table for details).
- Dynamic persistent binding.
- Unused devices in sd.conf no longer take up memory.
- Support for Dynamic Reconfiguration (DR). Driver version 6.x includes DR enhancements for SunFire 12K, 15K, E20K and E25K. (See the Compatibility section of this manual for prerequisites.)
- Improved scheduler algorithm addressing potential LUN starvation in specific environments.
- FDMI extension - Host Name support.

New Features in this Release

The Emulex Driver for Solaris version 6.02h:

- Support for 4 Gigabit/sec adapters (LP11000) with 1, 2, and 4 Gigabit auto negotiation.
- Solaris 10 support.
- HBAnyware utility is included as part of master kit: enabling GUI-based driver configuration and persistent binding management. (HBAnyware is not supported with Solaris version 2.6.)
- Different binding method on a per each HBA port basis.
- Resolved failover/failback issues.
- Resolved discovery issues with specific switches.
- Corrected queue depth for LP9802DC.

Prerequisites

- SPARC-based system.
- Solaris 2.6 (32-bit), Solaris 7 (32- and 64-bit), Solaris 8 (32- and 64-bit), Solaris 9 (32- and 64-bit) or Solaris 10 (32- and 64-bit).
- 64 MB system RAM.

Compatibility

- LP11002 and LP11000 (minimum firmware version 2.10a5)
- LPe11002 and LPe11000 (minimum firmware version 2.50a4)
- LP10000DC and LP10000 (minimum firmware version 1.90a4)
- LP9802DC (minimum firmware version 1.90a4)
- LP9802 (minimum firmware version 1.90a4)
- LP9402DC, LP9002DC, LP9002L and LP9000 (minimum firmware version 3.90a7)
- LP8000 and LP8000DC
 - If your HBA has a Dragonfly chip version 2.00 or greater, use firmware version 3.90a7.
 - If your HBA has a Dragonfly chip below version 2.00, use firmware version 3.30a7.

Refer to the LP8000 and LP8000DC Firmware Download page on the Emulex website to determine the Dragonfly chip version in use.

- LP7000E (minimum firmware version 3.21a0- this adapter does not support Dynamic Reconfiguration)
- Minimum OpenBoot PROM (OBP) levels:
 - Sun Fire 4800 OBP 5.15.2.
 - Sun Fire V880 OBP 4.7.5.
 - Enterprise 3500/4500 OBP 3.2.30.
 - Enterprise 250/450 OBP 3.26.0.
- Emulex recommends OpenBoot version 1.41a4 for all PCI HBAs for best DR compatibility.
- Emulex recommends OpenBoot version 2.41a5 for all SBus HBAs for best DR compatibility.

Things to Know Before You Download

- Before you download, create a temporary directory for the download package.

Caution: Installing version 6.02h of the driver for Solaris may be significantly different than in previous releases. Precisely follow the instructions for initial installation. Failure to follow these steps could render your system inoperable.

If you are updating the driver from version 6.00 or later to 6.02h in any system, refer to the following update procedure:

Updating a PCI Driver (from Version 6.x) on page 9.

If you are updating the driver from version 5.00i or later to 6.02h in any system, refer to the following update procedure:

Updating a PCI Driver (from Version 5.x) on page 10.

If you are updating the driver from version 4.21e or earlier to version 6.02h, refer to the appropriate update procedure:

Updating a PCI Driver (from Version 4.x) on page 11.

Updating an SBus Driver (from Version 4.x) on page 12.

Updating PCI and SBus Drivers in One System (from Version 4.x) on page 13.

After you install the driver, you must install the utility package, which includes lputil, HBAnyware and the HBA API libraries. The lputil utility is no longer included in the driver kit.

Installing Driver Utilities on page 15.

As a matter of course, you should always make a full system backup prior to beginning an installation or update procedure.

-
- If you are currently running an older version of the Emulex driver for Solaris, the update procedure includes removing the old driver.
 - Before upgrading the driver, you must uninstall any previously installed utilities. Refer to the "Uninstalling the Utilities" section on page 17 for instructions.
 - /etc/hostname.lpfc<0-N> - configuration files specifying adapters for use with IP networking. For each LightPulse HBA involved in networking, there must be a hostname.lpfc<0-N> file in /etc containing the host's name on that network.

Note: These files are not required if the driver is being used for SCSI support only.

Known Issues

The following issues have been reported at the time of publication. These issues may not yet have been verified or confirmed and may apply to another product, such as the driver or hardware.

FCP

Discovering Disk Drives: `drvconfig(1M)` and `disks(1M)`

Solaris provides a utility, `drvconfig(1M)`, which probes for new hardware added after the system is booted, first run `drvconfig`, then run `disks` to acquire access to new devices.

Sometimes, the Emulex driver for Solaris recognizes disk drives at boot time but you cannot use those drives because special file entries for them do not exist. In that case, run the `disks(1M)` command to create the special file entries in `/dev/dsk` and `/dev/rdisk`. In some cases, you may need to first run `drvconfig(1M)`, which creates the `/devices` tree, and then run `disks(1M)`. Since Solaris 8, `drvconfig(1M)` and `disks(1M)` are replaced by the command `devfsadm(1M)`.

Finally, in extreme cases, do the following, as 'root':

```
# touch /reconfigure
# reboot
```

A reconfiguration reboot makes Solaris rediscover all possible devices at boot time.

Different SCSI Target IDs on Different Hosts

A Fibre Channel target is assigned its `D_ID` at loop initialization time; the SCSI target ID for that target is assigned by the device driver when the device is first discovered. It is possible for the `D_ID` to change between one loop initialization and the next. Every time a system boots or a target is added to or removed from the Fibre Channel, the loop will be re-initialized. After a system has booted, it will maintain a constant view of the same target ID because the driver software remaps the SCSI target ID to the new `D_ID` on the fly. However, a second system may use a different SCSI target ID for that target. If you want to work with the same target across multiple hosts, you may find the same Fibre Channel target is known by different SCSI target IDs. Use the persistent binding feature to maintain a consistent target ID on all systems.

IP

Promiscuous Mode

The Emulex driver for Solaris supports promiscuous mode. However, promiscuous mode on Fibre Channel works differently than you might expect compared to promiscuous mode on Ethernet.

Promiscuous mode allows system administrators and others to do things such as run `tcpdump` and other utilities to find out what packets are being sent from or received by a node on a network. Unlike a true Ethernet driver, the Emulex driver working in promiscuous mode does not really receive all the packets going over the wire. Fibre Channel, unlike Ethernet, is not a broadcast medium. When promiscuous mode is enabled on the host machine, only the packets sent between the host machine and another machine can be seen. Packets sent between other machines cannot be seen. This restriction is a property of Fibre Channel itself.

An artifact of the driver's implementation also means that packets may not be reported in order. Sometimes a TCP acknowledgment packet can be sent back up a promiscuous STREAM before transmitted data packets are themselves sent up that STREAM.

Networking throughput decreases and latency increases for all packets sent or received through `lpfc` while promiscuous mode is enabled. Promiscuous mode forces a data copy for each incoming packet to make a duplicate copy for the socket operating in promiscuous mode. Another side effect of promiscuous mode is an increase in memory use.

Common Issues

Drive light turns out

Typically, the drive light does not turn on again by itself. The drive will come back on when a command to it times out and the driver tries to abort the command; or it may turn on at the next loop event. Contact the drive manufacturer for new firmware to fix the problem.

LIP type F8

Some disk drive firmware does not recover from a LIP type F8 in the way expected by other Fibre Channel devices; the result can be a loop full of devices that are all hopelessly confused. Reboot all the hosts involved and power cycle all the disk drives involved.

Unfortunately, it may be difficult to realize that an F8 LIP has happened: subsequent loop events may obscure the relevant log messages on the console or in the system log. The Emulex Digital Hub will bypass devices sending LIP F8, and keep the loop up and running.

Drive Firmware Update After Labeling Isn't Reported

If the drive firmware is updated after the disk is labeled by Solaris, the new firmware revision is not reflected in the Solaris disk label. Solaris stores the firmware revision number in the disk's ASCII name in the Solaris disk label; where it is displayed by the format command. If you update the drive's firmware, you must do an explicit inquiry command from within format to see the change, re-labeling doesn't help. Try re-formatting the disk.

Note: On Solaris/SPARC systems with the sd driver, disks labeled by Solaris/x86 aren't recognized and are reported as having "corrupt" labels; in that case, Solaris will issue an inquiry, so in this instance Solaris/SPARC will accurately report the updated firmware revision.

fsck Through Block Device Fails

Using fsck against the block- instead of the character-special device can result in error messages about unreadable blocks and then messages about partially allocated inodes, unknown file types and other unusual problems. Damage to the slice will result. Using the character-special device always works correctly.

Fibre Channel Network Doesn't Operate After pkgadd and reboot

Immediately after doing pkgadd and rebooting, lpfc isn't configured. An additional reboot is required to cure the problem. Sometimes, you may have to reboot a third time. Follow the installation instructions in this manual to avoid the problem.

SPARC Specific Issues

OBP Fails to Recognize LightPulse Correctly

SPARC systems with older Open Boot Prom revisions may fail to recognize the HBA. The lpfc driver will print out warning messages about being unable to configure the HBA. You must upgrade the OBP firmware or employ a workaround.

Installing the Emulex Driver for Solaris Version 6.02h

Introduction

Installing version 6.02h of the driver for Solaris may be significantly different from version 4.00 or 5.00. Precisely follow the instructions for initial installation. Failure to follow these steps could render your system inoperable.

The Emulex driver for Solaris download package, which you download from the Emulex Web site, includes a tar file for the driver and a tar file for the driver utilities (lputil, HBAnyware and HBA API files). When you install the download package, you will perform the following steps:

1. Perform the appropriate driver installation or upgrade procedure:
 - **Install a Driver for the First Time** on page 6.
 - If you are updating the driver from version 6.00 or later to 6.02h in any system, refer to the following update procedure:
 - **Update a PCI Driver (from Version 6.x)** on page 9.
 - If you are updating the driver from version 5.00i or later to 6.02h in any system, refer to the following update procedure:
 - **Update a PCI Driver (from Version 5.x)** on page 10.
 - If you are updating the driver from version 4.21e or earlier to version 6.02h, refer to the appropriate update procedure:
 - **Update a PCI Driver (from Version 4.x)** on page 11.
 - **Update an SBus Driver (from Version 4.x)** on page 12.
 - **Update PCI and SBus Drivers in One System (from Version 4.x)** on page 13.
2. Install the driver utilities.

Note: The HBAnyware and LightPulse (lputil) utilities are bundled together and must be installed separately from the driver. Refer to the "Installing the Utilities" section of this manual for more information.

Installing the Driver for the First Time

To install the Emulex driver for Solaris for the first time:

1. Login as or su to 'root'.
2. Load the package from your distribution medium into a directory, referred to here as <where-you-put-it>. The driver is a regular tar file, named lpfc-6.02h-sparc.tar.
3. Change to the <where-you-put-it> directory. Type:
`cd <where-you-put-it>`
4. Extract the installation image from the tar file. Type:
`tar xvf lpfc-6.02h-sparc.tar`
5. At the shell prompt, type:
`pkgadd -d `pwd``

If you are installing the lpfc driver on an alternate root, type:
`pkgadd R <alt-root> -d `pwd``

6. Answer the following question:

```
Select package(s) you wish to process (or 'all' to process all
packages). (default: all) [?,??,q]:
```

7. Select the number associated with lpfc for installation. If you need help, type? or ?? . You will see some additional questions related to directories. Ordinarily, the defaults are sufficient. To select a default press <Enter>.

- a. Answer the question:

```
Rebuild manual pages database for section 7d [y,n,?]:
```

Manual pages for lpfc are installed in section 7d of the online man pages. Normally the catman command can be run to create preformatted versions of the on-line manual from the nroff(1) input files. Each manual page is examined and those whose preformatted versions are missing or out of date are recreated. If any changes are made, catman recreates the windex database. Depending on your system, this operation can take anywhere from 1 to 10 minutes. If you type n , the catman will be skipped which will allow the installation to complete quickly. The catman command skipped will be output for the administrator to run at a later point in time, if desired.

- b. The next question you will see is:

```
Use IP networking over Fibre Channel [y,n,?]:
```

Answer Yes if you wish to enable IP or No if you do not wish to enable IP.

8. Continue with the installation by answering Yes to the following:

```
Do you want to continue with the installation of <lpfc> [y,n,?]:
```

9. The install package provides running commentary on the installation process. Be sure to examine the output for any errors or warnings. When the installation concludes, take notice of the message:

```
SCSI: If you are using lpfc to access disks, be sure to check the
configuration file of your SCSI target driver (presumably sd.conf)
to ensure that the driver will probe for all of the targets/luns
in your environment.
```

10. Finally you will be asked:

```
Select package(s) you wish to process (or 'all' to process all
packages). (default: all) [?,??,q]:
```

This is the same question that began your session. If you do not want to install additional packages, type q. The lpfc driver installation is now complete. After you quit, you will see the following message:

```
***IMPORTANT NOTICE ***
```

```
This machine must now be rebooted in order to ensure sane
operation.
```

At this point the driver is loaded and automapped targets are visible to the system, HBAnyware, and lputil.

11. Before you reboot the machine, consider whether you must modify any configuration files, beyond the defaults established by the installation process.

- a. Run HBAnyware to inspect the driver's configuration parameters or view the /kernel/drv/lpfc.conf file. (See the Configuration manual for details on starting and running HBAnyware.) If necessary, change settings as you deem appropriate. However, the defaults

are likely to be satisfactory. You can find `lpfc.conf` in the `/kernel/drv` if the driver was installed referencing the root directory.

- b. If you have FCP targets with multiple LUNs, add configuration entries for them to `/kernel/drv/sd.conf`. These files were set up by the installation process to probe for some FCP targets on the loop - but only for LUN 0 on some targets. If you have FCP devices with multiple LUNs or target IDs greater than 17, you must modify `sd.conf` to find them at boot time.
 - c. If you have more than one LightPulse adapter that will be configured for IP networking in the host, you must create an address file for each additional adapter. The installation process created `/etc/hostname.lpfc0` for the first adapter.
12. Reboot the system to ensure proper operation, if networking was enabled during installation or if non-dynamic driver parameters were changed. At the shell prompt, type the following:

```
# sync
# reboot
```

The system will reboot to multi-user mode and the Emulex LightPulse driver is available for use.

If you are enabling IP networking, you may find after rebooting the system in step 10 that FCP disk access works but host-to-host IP access does not work. The `lpfc` network interface isn't configured; as Solaris boots, it will print out an error message when configuring `lpfc0`, the Fibre Channel network, indicating that there is no such interface. Running `netstat -i` will also reveal that there is no `lpfc` interface. In this case, you must reboot the machine a second time. In certain rare cases, it may be necessary to reboot the machine a third time. Be careful, though: if `netstat -i` shows an `lpfc` interface, but the network still isn't being started, you probably have an error in your `hostname.lpfc<0-N>` file or some other error in your network configuration files.

Installing the Driver as a Loadable Kernel Module

When installing the device driver as a loadable kernel module, you must use supported versions of the hardware and software. Refer to the Compatibility section of the manual. Read the Configuration section before you start the installation.

To load the driver as a loadable module:

1. Login as 'root' or su to 'root'.
2. Type:

```
modload /kernel/drv/lpfc (for 32-bit platforms) or
modload /kernel/drv/sparcv9/lpfc (for 64-bit platforms)
```
3. Type:

```
cfgadm
```

Identify and record the `Ap_Id`'s of all the Emulex LightPulse adapters on the system.
4. Type:

```
cfgadm -c configure <Ap_Id>
```

for each of the Emulex LightPulse adapters.

To unload the driver:

1. Login as 'root'.
2. Un-mount all the filesystems that belong to the `lpfc` driver.
3. Un-plumb all IP interfaces that belong to the `lpfc` driver, if networking is enabled on the `lpfc` driver.

4. Stop all HBAnyware/rmsserver processes.
5. Type:
`cfgadm`
Record the Ap_Id's of all the Emulex LightPulse adapters on the system.
6. Type:
`cfgadm -c unconfigure <Ap_Id>`
for each of the Emulex LightPulse adapters.
7. Type:
`modunload -i <module_id>`
where <module_id> is the id of the lpfc driver. This id can be obtained by running modinfo.

Updating Existing Drivers

Use the following procedures to update an existing driver:

- Update a PCI Driver (from Version 6.X) on page 9.
- Update a PCI Driver (from Version 5.X) on page 10.
- Update a PCI Driver (from Version 4.X) on page 11.
- Update an SBus Driver (from Version 4.X) on page 12.
- Update PCI and SBus Drivers (from Version 4.X) in One System on page 13.

Updating a PCI Driver (from Version 6.X)

The driver can be updated by saving certain driver configuration files and using pkgm. Two files are saved when a pkgm is performed.

- The /kernel/drv/lpfc.conf file is moved to /usr/tmp/lpfc.conf.pkgm to allow you to restore any customized parameter settings after installing a new device driver.
- The file /kernel/drv/sd.conf is copied to /usr/tmp/sd.conf.pkgm. The file must be copied, because the sd driver still needs this file. Additionally any lpfc specific settings in the original file are deleted by the pkgm process.

After installing a new device driver, the saved files can be used to restore all the customized disk settings.

WARNING: The 6.02h driver adds a few new configuration parameters, deprecates a few, and renames a couple of old parameters. Do NOT replace the current lpfc.conf with lpfc.conf from a previous release.

Note: Before upgrading the driver, you must uninstall any previously installed utilities. Refer to the "Uninstalling the Utilities" section on page 17 for instructions.

To update the driver:

1. Make a full system backup.
2. Login as 'root', or su to 'root'.
3. Save the file /etc/path_to_inst to /usr/tmp/path_to_inst.lpfc. The /etc/path_to_inst file is not saved when a pkgm is performed.

4. At the shell prompt, type:
`pkgrm lpfc`
5. Copy the file `/usr/tmp/path_to_inst.lpfc` back to `/etc/path_to_inst`.
6. Install the new driver package. Type:
`pkgadd -d `pwd``
7. Restore any customized parameter settings in `sd.conf` and `lpfc.conf`:
 - Copy the `/usr/tmp/sd.conf.pkgrm` file back to `/kernel/drv/sd.conf`.
 - Update the configuration parameters in `/kernel/drv/lpfc.conf` with the values from `/usr/tmp/lpfc.conf.pkgrm`.
8. Reboot the system.

Updating a PCI Driver (from Version 5.X)

The driver can be updated by saving certain driver configuration files and using `pkgrm`. Two files are saved when a `pkgrm` is performed.

- The `/kernel/drv/lpfc.conf` file is moved to `/usr/tmp/lpfc.conf.pkgrm` to allow you to restore any customized parameter settings after installing a new device driver.
- The file `/kernel/drv/sd.conf` is copied to `/usr/tmp/sd.conf.pkgrm`. The file must be copied, because the `sd` driver still needs this file. Additionally any `lpfc` specific settings in the original file are deleted by the `pkgrm` process.

After installing a new device driver, the saved files can be used to restore all the customized disk settings.

WARNING: The 6.02h driver adds a few new configuration parameters, deprecates a few, and renames a couple of old parameters. Do NOT replace the current `lpfc.conf` with `lpfc.conf` from a previous release.

Note: Before upgrading the driver, you must uninstall any previously installed utilities. Refer to the "Uninstalling the Utilities" section on page 17 for instructions.

To update the driver:

1. Make a full system backup.
2. Login as 'root', or su to 'root'.
3. Save the file `/etc/path_to_inst` to `/usr/tmp/path_to_inst.lpfc`. The `/etc/path_to_inst` file is not saved when a `pkgrm` is performed.
4. At the shell prompt, type:
`pkgrm lpfc`
5. Copy the file `/usr/tmp/path_to_inst.lpfc` back to `/etc/path_to_inst`.
6. Install the new driver package. Type:
`pkgadd -d `pwd``
7. Restore any customized parameter settings in `sd.conf` and `lpfc.conf`:
 - Copy the `/usr/tmp/sd.conf.pkgrm` file back to `/kernel/drv/sd.conf`.
 - Update the configuration parameters in `/kernel/drv/lpfc.conf` with the values from `/usr/tmp/lpfc.conf.pkgrm`. This can be done without manual intervention by using the tool

/usr/sbin/lpfc/update_lpfc. Refer to “update_lpfc” on page 11 for more information on the tool.

8. Examine the file /tmp/lpfc.conf and make sure the tool has migrated the user settings from 5.x to 6.x correctly. Copy the file /tmp/lpfc.conf back to /kernel/drv/lpfc.conf
9. Reboot the system.

update_lpfc

update_lpfc is a simple conversion tool used to upgrade the Solaris lpfc driver from versions 5.01x or 5.02x to versions 6.00g, 6.01c, 6.01f, 6.02e, 6.02f, 6.02g or 6.02h. This tool converts the older lpfc.conf files to the proper 6.x syntax.

Usage:

```
/usr/sbin/lpfc/update_lpfc <lpfc.conf.5x> <lpfc.conf.6x>
```

The script extracts relevant information (user updates) from the original lpfc.conf (5x) and applies them to the new lpfc.conf (6x). The tool does not modify either of the input files. The output is directed to stdout. The tool employs sed and nawk to process the information from the lpfc.conf file.

Caveats

- The tool does not check for syntax in the original lpfc.conf file. The tool also does not validate the values of the configuration parameters, except for the ones described below. It is assumed that the original lpfc.conf is a working file.
- Comments from the original lpfc.conf are ignored.
- All configuration parameter assignments should start on column 1. There should be no leading white spaces in the configuration parameter assignments in the original lpfc.conf file.
- The user defined persistent bindings should not be interspersed with comments. For example, the following persistent bindings will not be translated to 6.x correctly.

```
fcplib-DID="0000ef:lpfc0t3",  
# "0000e8:lpfc0t4;  
"0000e4:lpfc0t4";
```
- The permissible values for the configuration parameter ‘automap’ have changed in the 6.x driver. The permissible values are 0 and 1. If ‘automap’ was set to 1, 2, or 3 in the original lpfc.conf file, then the tool will change this value to 1.
- The ‘fcplib-method’ parameter is new to lpfc 6.x driver. The tool sets ‘fcplib-method’ based on the type of persistent bindings used in the original lpfc.conf file. If persistent bindings are not defined in the original lpfc.conf file, then the tool sets ‘fcplib-method’ based on the automap parameter from the original lpfc.conf file.
- If ‘scan-down’ was set to 2 in the original lpfc.conf, then the tool will set ‘scan-down’ to 1. The 6.x lpfc driver does not allow a ‘scan-down’ value of 2.
- The deprecated lpfc.conf parameters are not carried over.
- If the configuration parameters ‘cr-count’, ‘num-bufs’ and ‘num-iocbs’ are set to their default settings in the original lpfc.conf, then the tool will update the values of these parameters to their new defaults.

Updating a PCI Driver (from Version 4.X)

The driver can be updated by saving certain driver configuration files and using pkgm. Two files are saved when a pkgm is performed:

- The /kernel/drv/lpfc.conf file is moved to /usr/tmp/lpfc.conf.pkgrm to allow you to restore any customized parameter settings after installing a new device driver.
- The /kernel/drv/sd.conf file is copied to /usr/tmp/sd.conf.pkgrm because the pkgrm process deletes any lpfc specific settings in the original file. After installing a new device driver, the saved file can be used to restore all the customized disk settings.

WARNING: The 6.02h driver adds a few new configuration parameters, deprecates a few, and renames a couple of old parameters. Do NOT replace the current lpfc.conf with lpfc.conf from a previous release.

Note: Before upgrading the driver, you must uninstall any previously installed utilities. Refer to the "Uninstalling the Utilities" section on page 17 for instructions.

To update the driver:

1. Make a full system backup.
2. Login as 'root', or su to 'root'.
3. Save the file /etc/path_to_inst to /usr/tmp/path_to_inst.lpfc. The /etc/path_to_inst file is not saved when a pkgrm is performed.
4. At the shell prompt, type:

```
pkgrm lpfc
```
5. Copy the file /usr/tmp/path_to_inst.lpfc back to /etc/path_to_inst.
6. Install the new driver package. Type:

```
pkgadd -d `pwd`
```
7. Restore any customized parameter settings in sd.conf and lpfc.conf:
 - Copy the /usr/tmp/sd.conf.pkgrm file back to /kernel/drv/sd.conf.
 - Update the configuration parameters in /kernel/drv/lpfc.conf with the values from /usr/tmp/lpfc.conf.pkgrm.
8. Reboot the system.

Updating an SBus Driver (from Version 4.X)

If you update from the 4.x SBUS driver, read through this section carefully. Otherwise, refer to the "Update a PCI Driver (from Version 5.X)" section on page 10.

The driver can be updated by saving certain driver configuration files and using pkgrm. Two files are saved when a pkgrm is performed.

- The /kernel/drv/lpfs.conf file is moved to /usr/tmp/lpfs.conf.pkgrm to allow you to restore any customized parameter settings after installing a new device driver.
- The /kernel/drv/sd.conf file is copied to /usr/tmp/sd.conf.pkgrm because the pkgrm process deletes any lpfs specific settings in the original file. After installing a new device driver, the saved file can be used to restore all the customized disk settings.

WARNING: The 6.02h driver adds a few new configuration parameters, deprecates a few, and renames a couple of old parameters. Do NOT replace the current lpfc.conf with lpfc.conf from a previous release.

Note: Before upgrading the driver, you must uninstall any previously installed utilities. Refer to the "Uninstalling the Utilities" section on page 17 for instructions.

To update the driver:

1. Make a full system backup.
2. Login as 'root', or su to 'root'.
3. Save the file /etc/path_to_inst to /usr/tmp/path_to_inst.lpfs. This file is not saved when a pkgrm is performed.
4. At the shell prompt, type:

```
pkgrm lpfs
```
5. Copy the file /usr/tmp/path_to_inst.lpfs back to /etc/path_to_inst.
6. Install the new driver package. Type:

```
pkgadd -d `pwd`
```

Caution: lpfs will be renamed to lpfc and the instance number of lpfs will be modified in the /etc/path_to_inst file.

The following messages for lpfs devices will be save into the file /usr/tmp/lpfs_new_instance when installation completed:

```
Move lpfs0 to lpfc0  
Move lpfs1 to lpfc1
```

Write down those messages, you will need them at the next step.

7. To restore any customized parameter settings in sd.conf and lpfc.conf, follow these instructions:
 - Copy the /usr/tmp/sd.conf.pkgrm file back to /kernel/drv/sd.conf.
 - Rename "lpfs" to "lpfc" in both /usr/tmp/lpfc.conf.pkgrm and /kernel/drv/sd.conf files.
 - Replace the old lpfs instance number to the new lpfc instance number in the /usr/tmp/lpfc.conf.pkgrm file.
 - Update the configuration parameters in /kernel/drv/lpfc.conf with the values from /usr/tmp/lpfc.conf.pkgrm.
8. Reboot the system.

Updating PCI and SBus Drivers (from Version 4.X) in One System

This procedure describes how to update Solaris driver version 4.x to version 6.02h on a system with both PCI and SBus Emulex HBAs installed.

The driver can be updated by saving certain driver configuration files and using pkgrm. Two files are saved when a pkgrm is performed:

- For 4.x PCI drivers, lpfc, the /kernel/drv/lpfc.conf file is moved to /usr/tmp/lpfc.conf.pkgrm to allow you to restore any customized parameter settings after installing a new device driver.
- For 4.x SBUS drivers, lpfs, the /kernel/drv/lpfs.conf file is moved to /usr/tmp/lpfs.conf.pkgrm.

- The file /kernel/drv/sd.conf is copied to /usr/tmp/sd.conf.pkgm because the pkgm process deletes any lpfc or lpfs specific settings in the original file depending on removing lpfc or lpfs. After installing a new device driver, the saved file can be used to restore all the customized disk settings.

WARNING: The 6.02h driver adds a few new configuration parameters, deprecates a few, and renames a couple of old parameters. Do NOT replace the current lpfc.conf with lpfc.conf from a previous release.

Note: Before upgrading the driver, you must uninstall any previously installed utilities. Refer to the "Uninstalling the Utilities" section on page 17 for instructions.

To update and merge the drivers:

1. Make a full system backup.
2. Login as 'root', or su to 'root'.
3. Save the file /etc/path_to_inst to /usr/tmp/path_to_inst.lpfc. The /etc/path_to_inst file is not saved when a pkgm is performed.
4. Save the file /kernel/drv/sd.conf to /usr/tmp/sd.conf.lpfc. This step is needed because the /kernel/drv/sd.conf.pkgm does not have the correct setting after step 4 and 5 are performed.
5. At the shell prompt, type:

```
pkgm lpfs
```
6. At the shell prompt, type:

```
pkgm lpfc
```
7. Copy the file /usr/tmp/path_to_inst.lpfc back to /etc/path_to_inst.
8. Install the new driver package. Type:

```
pkgadd -d `pwd`
```

Caution: lpfs will be renamed to lpfc and the instance number of lpfs will be modified in the file /etc/path_to_inst.

The following messages for lpfs devices will be save into the file /usr/tmp/lpfs_new_instance when installation completed.

```
Move lpfs0 to lpfc0  
Move lpfs1 to lpfc1
```

Write down those messages, you will need them at the next step.

9. To restore any customized parameter settings in sd.conf and lpfc.conf, follow these instructions:
 - a. Copy the /usr/tmp/sd.conf.lpfc file back to /kernel/drv/sd.conf.
 - b. Merge the files /usr/tmp/lpfs.conf.pkgm and /usr/tmp/lpfc.conf.pkgm into /usr/tmp/lpfc.conf.
 - c. Rename "lpfs" to "lpfc " in both /usr/tmp/lpfc.conf and /kernel/drv/sd.conf files.
 - d. Replace the old lpfs instance number to the new lpfc instance number in the /usr/tmp/lpfc.conf file.
 - e. Update the configuration parameters in /kernel/drv/lpfc.conf with the values from /usr/tmp/lpfc.conf.
10. Reboot the system.

Directory Structure

After installation, the following directories are created on the system.

Table 1: Directory Structure

| Directory | Description |
|---------------------|---|
| /usr/sbin/lpfc | Driver utilities (This directory is created after the utilities are installed.) |
| /usr/sbin/hbanyware | HBAnyware files (This directory is created after the utilities are installed.) |

Installing the Utilities

Follow these instructions to install HBAnyware and lputil on your system. For ease of installation, HBAnyware and lputil are bundled together.

Installing HBAnyware and lputil

Prerequisites

- lpfc driver 6.02h must be installed prior to installing the utilities.
- Java Runtime Environment:
Version 1.4 or later of the Java Runtime Environment (JRE) must be installed. HBAnyware will not run under earlier versions of the JRE.
The JRE and instructions for installation can be found at <http://java.sun.com/downloads/index.html>.

Procedure

To install HBAnyware and lputil from the tar file:

1. Copy the <HBAnyware_version>.tar.gz file to a directory on the install machine.
2. cd to the directory to which you copied the gz file.
3. Untar the file. Type:

```
gzcat <HBAnyware_version>.tar.gz | tar xvf -
```
4. su to 'root'.
5. At the shell prompt, type:

```
pkgadd -d `pwd`
```

Note: The utilities require the java runtime binaries and libraries, so their path must be included in the PATH environment variable. For example, if the java runtime binaries are in /usr/java/bin, then include this path in the PATH environment variable.

For example: (bash> export PATH="\$PATH:/usr/java/bin")

Installing the HBAnyware Security Configurator

Follow these instructions to install the HBAnyware Security Configurator on your system.

Prerequisites

- lpfcdriver 6.02h must be installed prior to installing the utilities.
- HBAnyware must be installed on the system.
- Java Runtime Environment:
Version 1.4 or later of the Java Runtime Environment (JRE) must be installed. lputil and HBAnyware will not run under earlier versions of the JRE.
The JRE and instructions for installation can be found at <http://java.sun.com/downloads/index.html>.

Procedure

To install the HBAnyware Security Configurator utility from a tar file:

1. Copy the <HBAnywareSSC_version>.tar.gz file to a directory on the install machine.
2. cd to the directory to which you copied the gz file.
3. Untar the file. Type:

```
gzcat <HBAnywareSSC_version>.tar.gz | tar xvf-
```
4. su to 'root'.
5. At the shell prompt, type:

```
pkgadd -d `pwd`
```

Removing an Existing Driver

The Emulex driver for Solaris can be removed by using pkgrm. All the lpfc files will be removed, including the file containing the driver itself. The next time the system is rebooted, the driver will not be loaded. All system resources, such as disk space and memory, will be reclaimed.

Two files are saved when a pkgrm is performed:

- The /kernel/drv/lpfc.conf file is moved to /usr/tmp/lpfc.conf.pkgrm. This file is used as a reference to restore any customized parameter settings after updating a new device driver.
- The /kernel/drv/sd.conf file is moved to /usr/tmp/sd.conf.pkgrm. This file is used as a reference to restore any customized parameter settings after updating a new device driver.

To remove the driver:

1. Login as 'root', or su to 'root'.
2. Remove HBAnyware. Type:

```
pkgrm HBAnyware
```
3. At the shell prompt, type:

```
pkgrm lpfc
```
4. Perform a reconfiguration reboot into single user mode by typing:

```
reboot -- -rs
```

5. Type the following command to clean up the old structures: (Solaris 7 or higher)

```
disks -C
```

6. Install the new lpfc device driver using the "Install a Driver for the First Time" instructions on page 6.

When you install the driver with pkgadd, two files are saved if they exist, just in case they are needed. The file /kernel/drv/lpfc.conf is saved in /usr/tmp/lpfc.conf.pkgadd. The file /kernel/drv/sd.conf is saved in /usr/tmp/lpfc.conf.pkgadd.

For further information on installing and removing drivers, consult the Solaris system administration documentation and the pkgadd(1M) and pkgrm(1M) manual pages.

Uninstalling the Utilities

Follow these instructions to uninstall the utilities.

Note: If the HBAnyware Security Configurator is installed, it must be uninstalled before uninstalling HBAnyware and lputil.

Uninstalling the HBAnyware Security Configurator

1. Log on as 'root'.

2. Type:

```
pkgrm HBAnywareSSC
```

Uninstalling HBAnyware and lputil

To uninstall HBAnyware and lputil:

1. Log on as 'root'.

2. Type:

```
pkgrm HBAnyware
```

Loading Manpages

Manual pages are installed in section 7d of the online man pages. Normally the 'catman' command can be run to create preformatted versions of the on-line manual from the nroff(1) input files. Each manual page is examined, and those whose preformatted versions are missing or out of date are re-created. If any changes are made, catman re-creates the windex database. Depending on your system, this operation can take anywhere from 1 to 10 minutes. If you press n, the catman will be skipped, which allows the installation to complete quickly. If the catman command is skipped, it can be run later.

SUN SAN Foundation Software Considerations

Emulex now has a separate FCA driver (emlxs) that supports the SUN SAN Foundation Software. SUN distributes and supports the Emulex emlxs driver. Since Emulex LightPulse HBAs are supported by both the lpfc and emlxs drivers, the driver installation scripts do not modify existing installations of the other driver type. For example, if emlxs is already installed on your system and you attempt to install the lpfc driver, then the lpfc driver will not claim any of the HBAs that the emlxs driver has already claimed. The driver will be installed properly on the system.

To change the driver HBA bindings, run the emlxdrv utility. The utility enables you to choose which driver (lpfc or emlxs) controls the Emulex adapters in the system. The utility is available on the Emulex web site as part of the FCA utilities package (<http://www.emulex.com/ts/downloads/solsfs/solsfs.html#>).

Configuration

Introduction

The Emulex driver for Solaris has many options that can be modified to provide for different behavior. You can change these options using the HBAnyware™ utility or the LightPulse® lputil utility.

Note: HBAnyware is not supported with Solaris version 2.6. Dynamic features that rely on HBAnyware are not available with Solaris version 2.6.

- The HBAnyware utility is a Java-based, user friendly graphical environment. Use HBAnyware to do any of the following:
 - Discover HBAs
 - Reset HBAs
 - Sort HBAs
 - Set up persistent binding
 - Set driver parameters
 - Update firmware on the local HBA or on remote HBAs
 - Update FC boot code (BootBIOS, OpenBoot or EFIBoot) on the local HBA or on remote HBAs

Note: HBAnyware's rmserver must be running on all remote hosts that are to be discovered and managed.

Remote capabilities of HBAnyware are subject to fabric zoning configuration. Remote hosts to be discovered and managed by HBAnyware must be in the same zone.

- The LightPulse utility (lputil) is a command-line interface. Use lputil to do any of the following:
 - Download PCI configuration data files
 - Discover HBAs
 - Reset HBAs
 - Update firmware on the local HBA
 - Update FC boot code on the local HBA
 - Enable the BootBIOS message

Starting HBAnyware for Solaris

To start HBAnyware for Solaris:

1. su to 'root'.
2. Run the script:
`/usr/sbin/hbanyware/hbanyware`

To start the HBAnyware Security Configurator for Solaris:

1. su to 'root'.
2. Run the script:
`/usr/sbin/hbanyware/ssc`

Starting HBAnyware in Local Element Manager Mode

The HBAnyware utility can also be launched with a command line call for Solaris systems.

To launch the HBAnyware utility from the command line:

1. Type "HBAnyware" and press <ENTER>. This starts HBAnyware running using in-band access. You can also start the utility running in out-of-band access by adding an argument in the form "h=<host>". The <host> argument may be either the IP address of the host or its system name. The call will use a default IP port of 23333, but you can override this by optionally appending a colon (:) and the IP port.

Note: Remember that not all HBAs for a specific host may be running in-band. Therefore, running that host out-of-band may display HBAs that do not appear when the host is running in-band.

Examples of Modifications

- HBAnyware h=138.239.82.2
HBAnyware will show HBAs in the host with the IP address 138.239.82.2.
- HBAnyware h=Util01
HBAnyware will show HBAs in the host named Util01.
- HBAnyware h=138.239.82.2:4295
HBAnyware will show HBAs in the host with the IP address 138.239.82.2 using IP Port 4295.
- HBAnyware h=Util01:4295
HBAnyware will show HBAs in the host named Util01 using IP port 4295.

Run this modified command line to launch the HBAnyware utility for a single, remote host in local mode.

HBAnyware Window Element Definitions

The **HBAnyware** window contains five basic components: the menu bar, the toolbar, the discovery tree, the property tabs and the status bar.

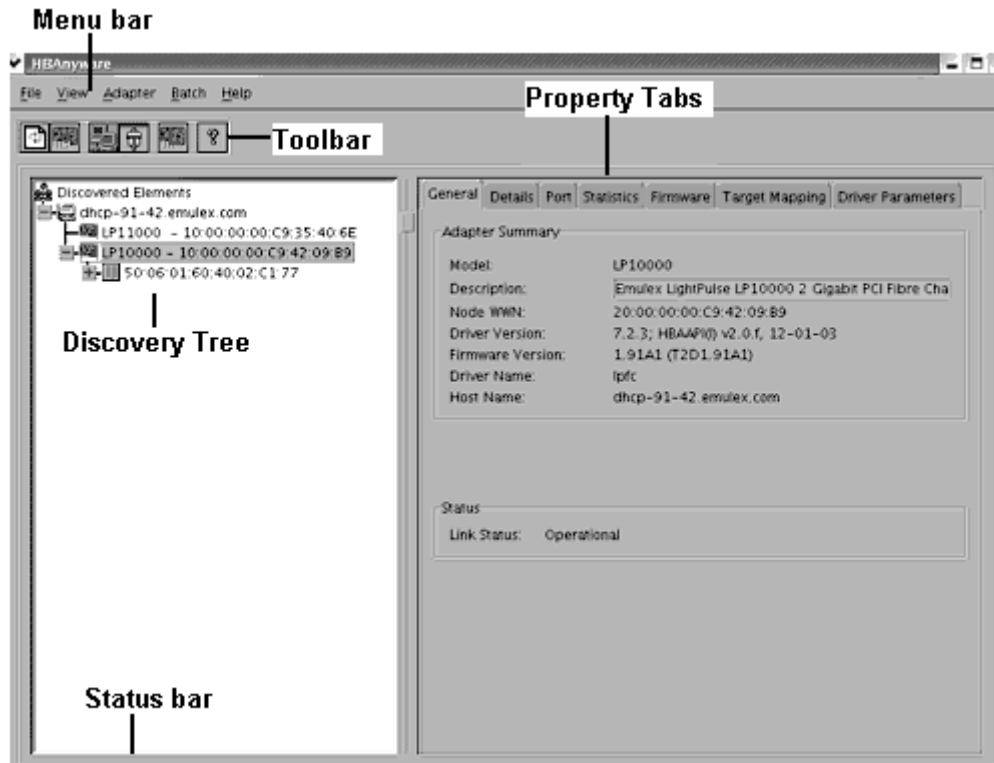


Figure 1: The HBAnyware Window with Element Call Outs

Note: The element you select in the discovery tree determines whether a menu item or toolbar icon is active. For example, if you select the local host or other system host, the Reset Adapter item on the Adapter menu is unavailable. The Reset Adapter toolbar button is unavailable as well.

The Menu Bar

The menu bar contains command menus that enable you to perform a variety of tasks such as resetting host bus adapters and sorting items in the discovery tree view. Many of the menu bar commands are also available from the toolbar.

The Toolbar

The toolbar contains buttons that enable you to refresh the discovery tree view, reset the selected host bus adapter and sort the discovery tree view. Many of the toolbar functions are also available from the menu bar.



Figure 2: The HBAnyware Toolbar

The toolbar is visible by default. Use the Toolbar item in the View menu to hide the toolbar. If the item is checked, the toolbar is visible.

The Toolbar Icon Definitions

The toolbar buttons perform the following tasks:



Click the **Rediscover** button to refresh the discovery tree display.



Click the **Reset** button to reset the selected host bus adapter.

Sort and Display Icons

Discovered adapters can be sorted by host name or fabric addresses. You can also choose to display only local or remote HBAs. See page 34 for details on sorting.



Group HBAs by Host Name (default)



Group HBAs by Fabric Address



Local HBAs Only



Help

The Discovery Tree

The discovery tree (left pane) displays icons representing discovered network (SAN) elements (local host name, system host names and all host bus adapters that are active on each host). Targets and LUNs, when present, are also displayed.

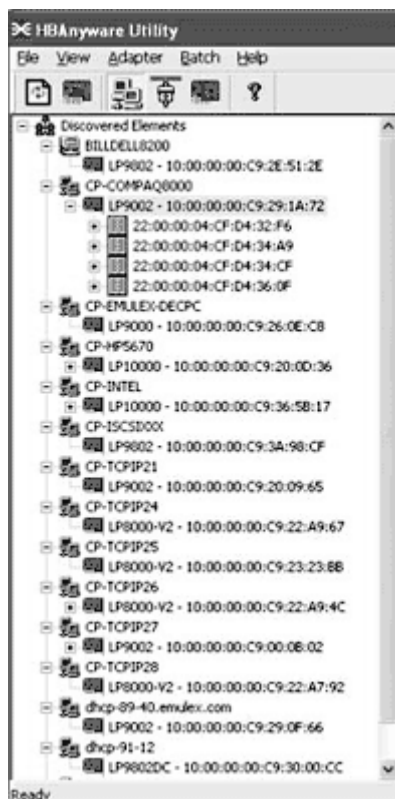


Figure 3: The HBAAnyware Discovery Tree

Discovery Tree Icons

Discovery tree icons represent the following:



The local host.



Other Hosts connected to the system.



A green HBA icon with black descriptive text represents an HBA that is online.



A green HBA icon with red descriptive text represents an HBA that is offline or otherwise temporarily inaccessible. Several situations could cause the HBA to be offline or inaccessible:

- The HBA on a local host is not connected to the SAN (Storage Area Network), but is still available for local access.
- The HBA on a local host is malfunctioning and is inaccessible to the local host as well as to the SAN.

- The HBA on a local host is busy performing a local firmware download and is temporarily inaccessible to the local host as well as to the SAN.



The Target icon represents connections to individual storage devices.



LUN 1 The LUN icon represents connections to individual LUNs.

Property Tabs

The property tabs display configuration, statistical and status information for network elements. The set of available tabs is context-sensitive, depending on the type of network element or HBA currently selected in the discovery tree.

Status Bar

The status bar is visible by default. Use the Status Bar item in the View menu to hide the status bar. If the item is checked, the status bar is visible.

Using the HBAnyware Command-Line Interface

The CLI (command-line interface) Client component of HBAnyware provides access to the capabilities of the Remote Management library from a console command prompt. This component is intended for use in scripted operations from within shell scripts, batch files, or the specific platform equivalent.

Note: HBAnyware must be running on all remote hosts that are to be discovered and managed. Remote capabilities of HBAnyware are subject to fabric zoning configuration. Remote hosts to be discovered and managed by HBAnyware must be in the same zone.

The CLI Client

The CLI Client is a console application named HBACMD. Each time you run this application from the command line, a single operation is performed.

The first parameter of this command is the requested operation. When the specified operation is completed, the command prompt is displayed. Most operations retrieve information about an entity on the SAN and display that information on the console.

Most of the CLI Client commands require one or more additional parameters that specify the nature of the command. A parameter used by many HBACMD commands specifies the World Wide Port Name of the host bus adapter (HBA) that is the target of the command. For example, the following command displays the port attributes for the HBA with the specified World Wide Port Name:

```
/usr/sbin/hbanyware/hbacmd portattrib 10:00:00:00:c9:20:20:20
```

Entering `/usr/sbin/hbanyware/hbacmd <no qualifiers>` displays a list of command options.

CLI Client Command Reference

Version

Syntax: `./hbacmd VERSION` (This command is not case sensitive.)

Description: The current version of the HBAnyware CLI Client application

Parameters: N/A

List HBAs

Syntax: `./hbacmd LISTHBAS` (This command is not case sensitive.)

Description: A list of the discovered manageable Emulex HBAs and their World Wide Node Names.

Parameters: N/A

Display HBA Attributes

Syntax: `./hbacmd HBAAttrib <wwpn>`

Description: A list of attributes for the HBA with the specified World Wide Port Name.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Port Attributes

Syntax: `./hbacmd PortAttrib <wwpn>`

Description: A list of attributes for the port with the specified World Wide Port Name.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Port Statistics

Syntax: `./hbacmd PortStat <wwpn>`

Description: A list of statistics for the port with the specified World Wide Port Name.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Server Attributes

Syntax: `./hbacmd ServerAttrib <wwpn>`

Description: A list of attributes for the specified server.

Parameters: `wwpn`- The World Wide Port Name of the server.

Download

Syntax: `./hbacmd DOWNLOAD <wwpn> <filename>`

Description: Loads the specified firmware image to the (HBA) with the specified WWPN.

Parameters: `wwpn`- The World Wide Port Name of the HBA that is the target of the firmware download.
The HBA can be either local or remote.

Filename- The pathname of the firmware image that is to be loaded. This can be any file that is accessible to the CLI client application, but we recommend that you keep image files in the Emulex Repository folder or directory.

Reset Adapter

Syntax: `./hbacmd RESET <wwpn>`

Description: Resets the HBA with the specified World Wide Port Name.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Target Mapping

Syntax: `./hbacmd TargetMapping <wwpn>`

Description: A list of mapped targets for the port with the specified World Wide Port Name.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Persistent Binding

Syntax: `./hbacmd PersistentBinding <wwpn> <source>`

Description: This function returns a list of the current persistent binding data associated with the HBA specified by ObjectPort. The data may be retrieved either from the driver itself (live), or from a configuration file.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Source- Either C or L. C = Configuration. L = Live.

All Node Info

Syntax: `./hbacmd AllNodeInfo <wwpn>`

Description: This functions retrieves target information for all targets that are visible to the specified HBA. This includes all automapped, persistently-bound, and unmapped targets. Because this function returns information for any unmapped targets, it is a more inclusive call than Persistent Binding.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Set Persistent Binding

Syntax: `./hbacmd SetPersistentBinding <wwpn> <scope> <bindtype> <id> <scsibus> <scsitarget>`

Description: Creates a persistent binding between an FCP target and OS SCSI information.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

Scope- P, I, or B. P = Bind set permanently. I = Bind set immediately. B = Bind set immediately and permanently at reboot.

BindType- D, P, or N. D = Enable binding by D_ID. P = Enable binding by WWPN. N = Enable binding by WWNN. ID- Either WWPN, WWNN, or D_ID (depending on BindType).

Remove All Persistent Binding

Syntax: `./hbacmd RemoveAllPersistentBinding <wwpn>`

Description: Removes all persistent bindings for the specified HBA. Only the configured bindings can be removed; rebooting is required to remove a live bindings.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

Remove Persistent Binding

Syntax: `./hbacmd RemovePersistentBinding <wwpn> <bindtype> <id> <scsibus> <scsitarget>`

Description: Removes a selected persistent binding. Only the configured bindings can be removed; rebooting is required in order to remove a live binding.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

BindType- D, P, or N. D = Enable binding by D_ID. P = Enable binding by WWPN. N = Enable binding by WWNN. ID- Either WWPN, WWNN, or D_ID (depending on BindType).

Binding Capabilities

Syntax: `./hbacmd BindingCapabilities <wwpn>`

Description: The flags returned by this function represent all binding capabilities present in the HBA specified by `ObjectPort`.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

Binding Support

Syntax: `./hbacmd BindingSupport <wwpn> <source>`

Description: This function returns the subset of capabilities that is currently active on the specified HBA.

Parameters: `wwpn`- The World Wide Port Name of the HBA. The HBA can be either local or remote.

Source- Either C or L. C = Configuration. L = Live.

Set Binding Support

Syntax: `./hbacmd SetBindingSupport <wwpn> <bindflag>`

Description: This function installs a set of active capabilities in the specified HBA.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

Bindflag- D, P, N, A, DA, Pa, or NA. D = Enable binding by D_ID. P = Enable binding by WWPN. N = Enable binding by WWNN. A = Enable binding by AUTOMAP. DA = Enable binding by D_ID and AUTOMAP. PA = Enable binding by WWPN and AUTOMAP. NA = Enable binding by WWNN and AUTOMAP.

Driver Parameters

Syntax: `./hbacmd DriverParams <wwpn>`

Description: This function returns the driver parameters array of the specified HBA. Each entry in the array contains the parameter name and values for minimum value, maximum value, current value, and default value.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

Set Driver Parameters

Syntax: `./hbacmd SetDriverParams <wwpn> <ctrlword> <param> <value>`

Description: This function is used to assign a value to a member of the Driver Parameters array belonging to the HBA referenced by `ObjectPort`. Only one parameter can be set for each call to this function.

Parameters: `wwpn`- The World Wide Port Name of the port. The port can be either local or remote.

ctrlword- P, G, B or N. P = Permanent. G = Global. B = Both. N = Neither

Set BootBIOS

Syntax: `./hbaCmd SetBootBios <wwpn> <ctrlword>`

Description: This function is used to enable/disable a BootBIOS firmware file that is present on an HBA. When you download a firmware file which has a Boot BIOS file attached, you have an option to enable or disable this boot file, depending upon the current state of this boot file.

Parameters: ctrlword- E or D. E = Enabled. D = Disabled.

Starting lputil

The LightPulse Diagnostic utility (lputil) is installed automatically when the driver utilities kit is installed.

Start the utility by entering the complete path to lputil. The path reflects the default installation path. If the installation path changed, you will need to adjust the command appropriately.

To start lputil type:

```
/usr/sbin/lpfc/lputil
```

No-Reboot Features

The Emulex 6.02h SPARC driver enables you to do the following without rebooting:

Note: HBAnyware is not supported for Solaris version 2.6.

- Discover newly provisioned Targets and LUNs.
- Add new adapters to the system.
- Make persistent binding changes.
- Make other driver configuration changes.
- Turn on the log capability (log-verbose).
- Update adapter firmware.
- Load or unload the driver.

Discovering New Targets and LUNS

For Solaris 9 and 10 Users:

Perform these steps to make targets and LUNs visible even if the unassigned targets and LUNs weren't contained in sd.conf since the last reboot.

To discover new targets and LUNs:

1. Edit the sd.conf file. (Ensure that the new targets and LUNs are identified in this configuration file.)
2. Run "update_drv -f sd". Newly provisioned targets and LUNs should now be seen.

Note: Note the space between -f, for "force", and sd, for the conf file to be re-read

See the man pages for more information on this command.

3. If LUNs are not available, run "devfsadm" from the command prompt.

For Solaris 8 Users- Unassigned Targets and LUNs Provisioned in sd.conf at Last Reboot

Perform the following steps to make targets and LUNS visible to the system:

- If the Automap parameter of lpfc.conf was turned ON at the last reboot, use the "devfsadm" command to make targets and LUNS visible to the system. See the man pages for more information on this command.
- If the Automap parameter was turned OFF at the last reboot, you can specify the SCSI ID using HBAnyware. These bindings take effect immediately without unloading and reloading the driver.

Note: The Emulex driver for Solaris does not take up memory space for unassigned targets.

Solaris 2.6, 7 or 8 Users - Unassigned Targets and LUNs not provisioned in sd.conf

This procedure enables new Targets and LUNs to be seen when the unassigned Targets and LUNs were not previously specified in sd.conf.

Note: On some SPARC systems (i.e. E250 and E450), the boot drive uses the sd driver to boot. You cannot unload and reload sd on these systems and therefore you must reboot unless you are using Solaris 9 or 10.

To unload sd-attached applications and unload/reload the sd driver:

1. Ensure the automap parameter in lpfc.conf is turned on (value =1) since the last reboot.
2. To unload the sd driver, all other drivers and applications attached to it must be unloaded or halted. In cases where the boot driver is controlled by a driver other than sd (i.e. E3/4/5/6500, V880, SunFire 4800), LUNs can be added dynamically with a few simple steps.
3. Unmount file systems attached to sd driver. Type:

```
umountall
```
4. Stop the vold daemon. Type:

```
# /etc/rc2.d/S92volmgt stop
```
5. Identify the module ID of sd driver and unload it. Type:

```
# modinfo | grep sd
```

A SunFire V880 displays:
19 101efd55 10a60 118 1 ssd (SCSI SSA/FCAL Disk Driver 1.147)
138 78148000 183d0 32 1 sd (SCSI Disk Driver 1.359)

In this case type:

```
# modunload -i 138
```

(The module ID is from the message above.)
6. Edit /kernel/drv/sd.conf to add or remove LUNs and/or targets.
7. Run the commands to discover new LUNs and/or targets. Type:

```
# devfsadm
```

This command automatically loads the sd driver, scans for new devices, and creates the necessary /dev/dsk and /dev/rdisk links.

If the devfsadm command is not recognized type:

```
# drvconfig  
# disks
```
8. Re-start the vold daemon. Type:

```
# /etc/rc2.d/S92volmgt start
```

9. Mount the file systems. Type:

```
# mountall
```

Removing Ghost LUNS

Ghost devices are targets that have been disconnected since the last reboot, but remain in the Solaris sd configuration. While the presence of these ghost devices does not create any runtime issues, they can slow down certain processes and also cause confusion.

The procedure for removing ghost devices is identical to the one used for discovering targets and LUNs. Refer to the "Discovering New LUNS and Targets" topic on page 27 to learn how to remove ghost devices.

Adding New Adapters, Changing Bindings and Changing Driver Parameters

For Solaris 9 and 10 Users

To add new adapters, change bindings and change driver parameters:

1. Login as 'root' or su to 'root'.
2. If you are adding an adapter not previously seen by the server:
 - a. Check for the following three conditions:
 - Ensure your Solaris host supports DR.
 - Use adapter firmware that supports DR.
 - Ensure the adapters are installed in DR capable PCI slots.
 - b. Power off a DR enabled PCI slot.
 - c. Insert the adapter and power up the slot.
3. Make driver parameter changes using HBAnyware. (Refer to the HBAnyware topics in this manual for more information.) Be sure to make changes Static. Static changes make the changes in memory and modify the lpfc.conf file so that if the host is rebooted, the driver parameters retain their intended values.

Note: Some driver parameters are dynamic. If a reboot is necessary, HBAnyware will prompt you to reboot. In that case, complete steps #4 - #9. If you had to unconfigure the adapters in step 7, you'll have to reconfigure them in step 11.

4. Stop the rmserver. (Stop HBAnyware.)
5. If you need to change or add target numbers and/or LUN numbers in sd.conf, make these modifications now.
6. Run the "cfgadm" command. You will see a list of controllers and their Ap_id numbers.
7. Unconfigure all lpfc driver instances using their Ap_id. Type: "cfgadm -d unconfigure Ap_id ..."
8. Type "modinfo | grep lpfc" and note the process number at the beginning of the returned string. It may look similar to: 85 7ba44000 525d0 269 1 lpfc (Emulex Lightpulse FC SCSI/IP).
9. Type "modunload -i 85". Replace "85" with your host's module-id number.
10. Run "update_drv -f sd". This forces Solaris to re-read sd.conf.
11. Run "cfgadm -c configure Ap_id Ap_id Ap_id". This example shows how to configure the entire list of lpfc driver instances you unconfigured in step 7 using one command. If you unconfigured more than one lpfc instance, separate each Ap_id with a space.

For Solaris 8 Users

The procedure for Solaris 8 and earlier is a little different because you may not be able to unload the sd driver on the fly. This procedure does not rely on DR support.

To add new adapters, change bindings and change driver parameters:

1. Login as 'root' or su to 'root'.
2. If you booted the system via scsi, you cannot use this procedure as root is linked to the scsi drive. Proceed to step 3 and then to step 13.
3. Make driver parameter changes using HBAnyware. (Refer to the HBAnyware topics in this manual for more information.) Be sure to make changes Static. Static changes make the changes in memory and modify the lpfc.conf file so that if the host is rebooted, the driver parameters retain their intended values.

Note: Some driver parameters are dynamic. If a reboot is necessary, HBAnyware will indicate this. In that case, complete steps #4 - #12.

4. If you need to change or add target numbers and/or lun numbers in sd.conf, make these modifications now.
5. Type "modinfo | grep sd" and note the module-id number at the beginning of the returned string. It may look similar to: 121 8cb16704 316b1 373 1 sd
6. Type "modunload -i 121". Replace "121" with the host's process number.
 - a. If this fails, stop the vold daemon. Type: "# /etc/rc2.d/S92volmgt stop".
 - b. If this is successful, again attempt to modunload the sd driver. If this fails, see step 11.
7. If the sd driver unloaded successfully, stop the rmserver. (Stop HBAnyware.)
8. Type "modinfo | grep lpfc" and note the process number at the beginning of the returned string. It may look similar to: 85 7ba44000 525d0 269 1 lpfc (Emulex Lightpulse FC SCSI/IP).
9. Type "modunload -i 85". Replace "85" with the host's process number.
10. Type "devfsadm". The sd and lpfc drivers will be reloaded automatically. The Volume Management daemon can now be started manually (i.e # /etc/rc2.d/S92volmgt start). If new devices do not appear continue to step 11.
11. Verify one last time that the lpfc.conf parameters contain the values you want and reboot the host to enable them.

Turning on Logging Capability

The Solaris 6.02h driver enables you to turn on or off the log-verbose parameter for problem tracking using HBAnyware. Refer to the HBAnyware topics in this manual for more information. (This capability is not supported for Solaris version 2.6.)

No-Reboot Firmware Updates

Emulex is the only vendor providing dynamic adapter firmware update during operation, and without stopping I/O traffic. You can dynamically update host bus adapter firmware during operation, and without stopping I/O traffic, using HBAnyware. Refer to the "Update Firmware" topics in this manual for more information about using HBAnyware. (This capability is not supported for Solaris version 2.6.)

Loading or Unloading the Driver Without Rebooting

Note: When the Solaris operating system is installed on a Fibre Channel drive, you must reboot the system because you cannot quiesce all I/O on the OS drive.

Systems must support dynamic reconfiguration.

To load the driver without rebooting:

1. Load the driver using the “modload” command.
2. Use the “cfgadm” command to configure Emulex HBAs.
3. Restart I/O.

To unload the driver without rebooting:

1. Quiesce all I/O on the device.
2. Use the “cfgadm” command to disconnect Emulex HBAs.
3. Unload the driver using the modunload command.

HBA Tasks

Discovering HBAs

You can discover adapters using either HBAnyware or lputil.

- HBAnyware allows you to discover both local and remote adapters.
- lputil allows you to discover local adapters only.

Discovering HBAs Using HBAnyware

Local and remote host bus adapters (HBAs) are discovered automatically when you launch HBAnyware.

Note: HBAnyware must be installed and the rmserver process(es) must be running on all remote hosts that are to be discovered and managed.

Remote capabilities of HBAnyware are subject to fabric zoning configuration. Remote hosts to be discovered and managed by HBAnyware must be in the same zone.

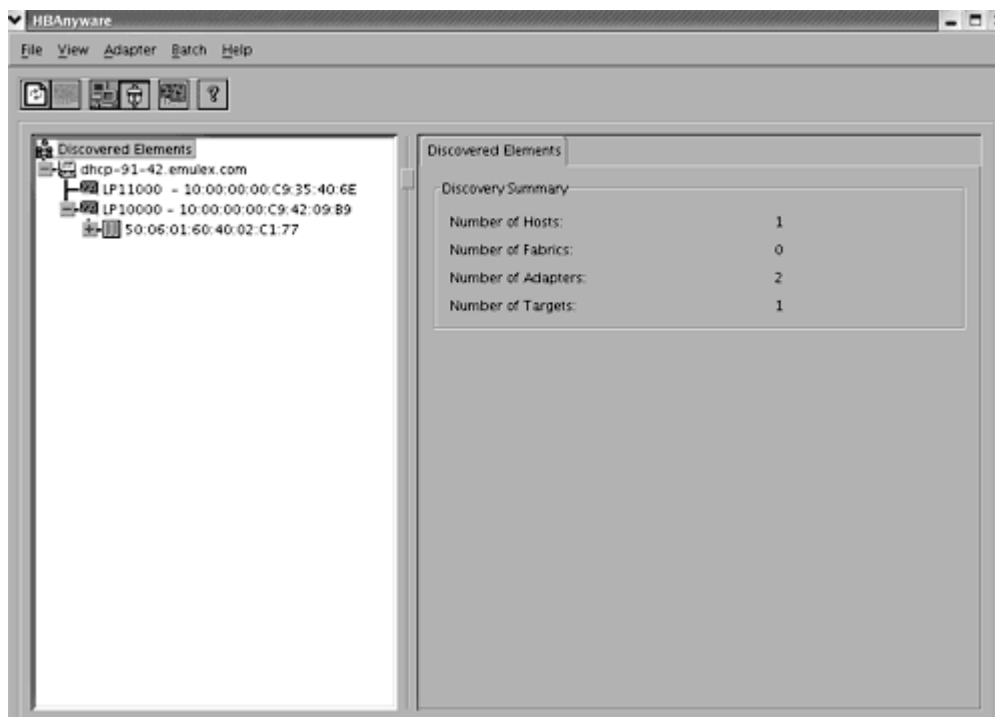


Figure 4: The HBAnyware Discovered Elements Tab

Configuring Discovery Settings

Using the **HBAnyware Discovery Settings** dialog box, you can configure the way HBAnyware performs discoveries. For example, you can set when the discovery server starts, the amount of time discoveries are refreshed, and when to remove previously discovered HBAs that are no longer being discovered.

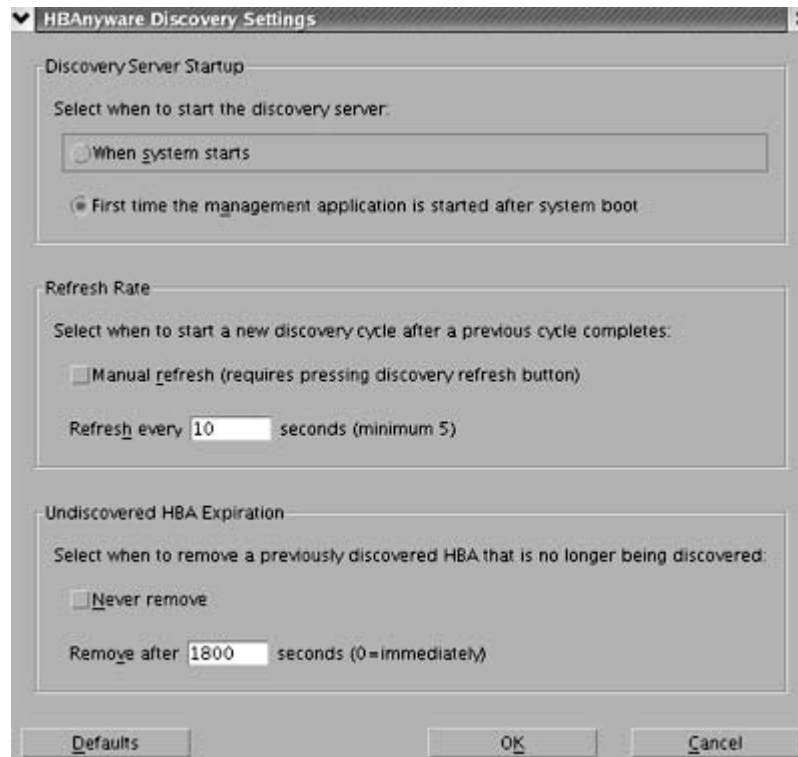


Figure 5: The HBAnyware Discovery Settings dialog box

To configure HBAnyware discovery settings:

1. From the Menu bar, click **View** then click **Discovery Settings**. The **HBAnyware Discovery Settings** dialog box appears.
2. Edit the settings and click **OK**.

Click **Default** to return the dialog box to its default settings.

Discovering HBAs Using lputil

When you start the LightPulse utility (lputil), all adapters are discovered and listed with information such as the host adapter number, instance number (i.e. lpfc0), board model type and whether the adapter is ready to use.

From the Main menu, enter 1, List Adapters.

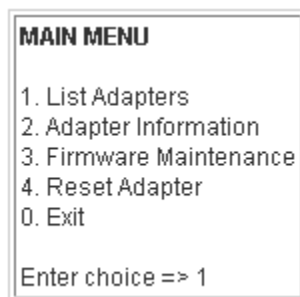


Figure 6: The lputil Main Menu


Sorting HBA Information

Sorting HBAs Using HBAnyware

You can use HBAnyware to sort the way discovered HBAs are displayed. You can sort HBAs by host name or fabric name.

You can also choose to view local HBAs or remote HBAs. By default, both local and remote HBAs are displayed by host name/fabric name.

- Switch between host name or fabric ID in one of two ways:
 - From the menu bar: click **View**, then click **Group HBAs by Host Name, Group HBAs by Fabric Address**. The current adapter display mode is checked.
 - From the toolbar, click one of the following buttons:

Group HBAs by Host Name (default). 

Group HBAs by Fabric Address. 

- HBAnyware sorts in ascending order. The sort recognizes letters, numbers, spaces and punctuation marks.

Group HBAs by Host Name

- Initially sorts by host name. Host names cannot be changed using HBAnyware; names must be changed locally on that system.
- Within each host system, sorts by HBA model.
- If multiple HBAs have the same model number, sorts models by world wide node name.
- If targets are present, sorts by world wide port name. Multiple HBAs may refer to the same target.
- If LUNs are present, sorts by LUN name.


Group HBAs by Fabric Address

- Initially sorts by fabric ID.
- Within each fabric ID, sorts by HBA model.

- If multiple HBAs have the same model number, sorts models by world wide node name.
- If targets are present, sorts by world wide port name. Multiple HBAs may refer to the same target.
- If LUNs are present, sorts by LUN name.
- If the fabric ID is all zeros, no fabric attachment is present.

Sorting Local HBAs Only Using HBAnyware

To display local HBAs only, do one of the following:

- From the menu bar: click **View**, then click **Local HBAs Only**. The current adapter display mode is checked.
- From the toolbar, click the  button.

Sorting Local HBAs Using Iputil

Local HBAs are automatically displayed.

Viewing HBA Information Using HBAnyware

Viewing Discovered Elements

The **Discovered Elements** tab in HBAnyware contains a general summary of the discovered elements. The Discovered Elements node is the root of the discovery tree, but it does not represent a specific network element. Expanding it will reveal all hosts, LUNs, targets and adapters that are visible on the SAN.

To view the discovered elements:

1. Start HBAnyware.

2. Click **Discovered Elements** in the discovery tree.

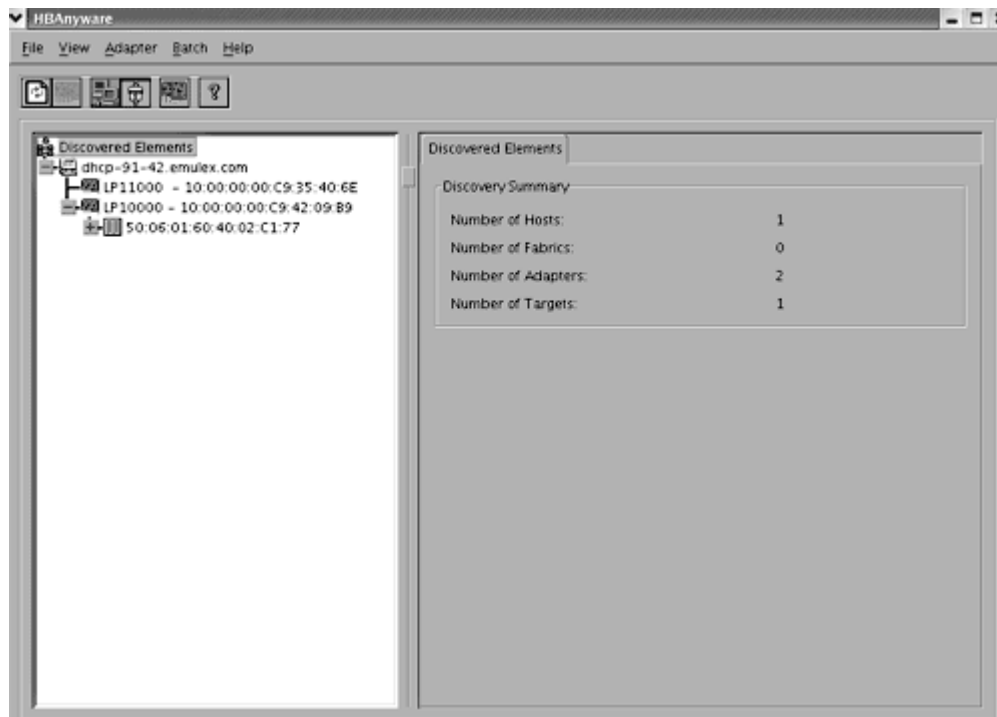


Figure 7: The HBAware Discovered Elements Tab

Discovered Elements Field Definitions

Number of Fabrics - the total number of fabrics discovered.

Number of Hosts - the total number of host computers discovered. This includes servers, workstations, personal computers, multiprocessors and clustered computer complexes.


Number of Adapters - the total number of host bus adapters (HBAs) discovered.

Number of Targets - the total number of unique targets discovered on the SAN. In the discovery tree, the same target can appear under more than one HBA.

Viewing Host Attributes

The **Host Attributes** tab in HBAware contains information specific to the selected host.

To view the host attributes:

1. Start HBAware.
2. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Host Name**.
 - From the toolbar, click the  button.

3. Click a host name in the discovery tree.

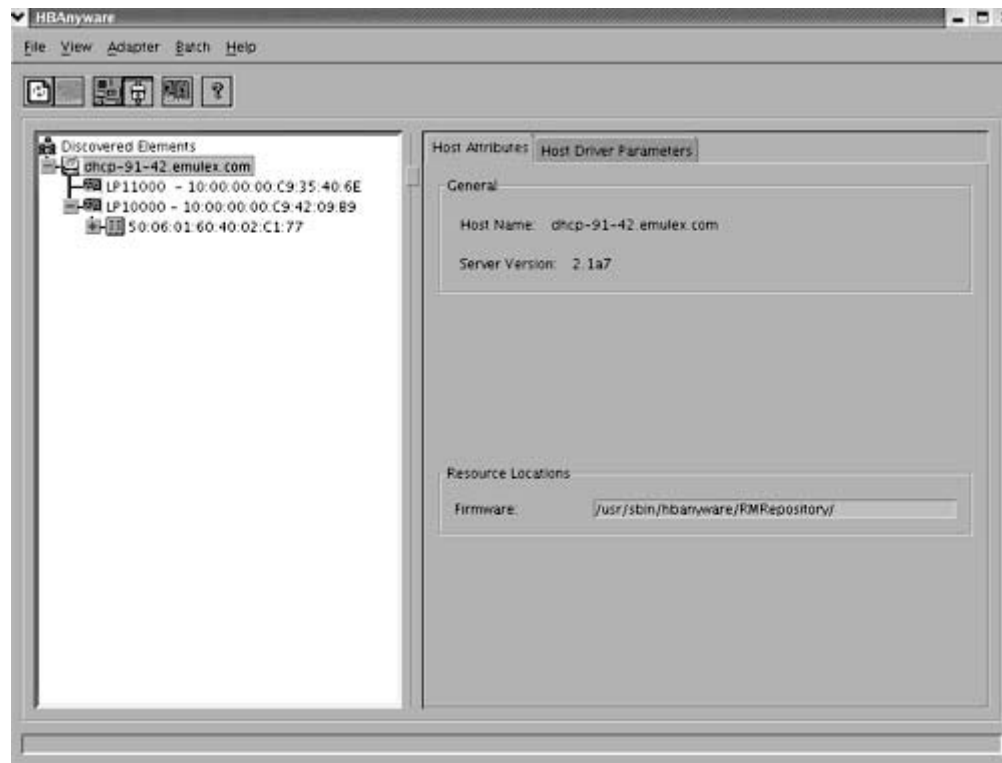


Figure 8: The HBAAnyware Host Attributes Tab

General Area Field Definitions

Host Name - the name of the host.

Server Version - the version number of the utility in use.


Resource Location Field Definitions

Firmware - the directory path where the firmware image files are moved prior to being downloaded to the HBAs on that host.

Viewing Target Attributes

The **Target Attributes** tab in HBAAnyware contains information specific to the selected target.

To view target attributes:

1. Start HBAAnyware.
2. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Host Name**.
 - From the toolbar, click the **Group HBAs by Host Name** button. 

3. Click a target in the discovery tree.

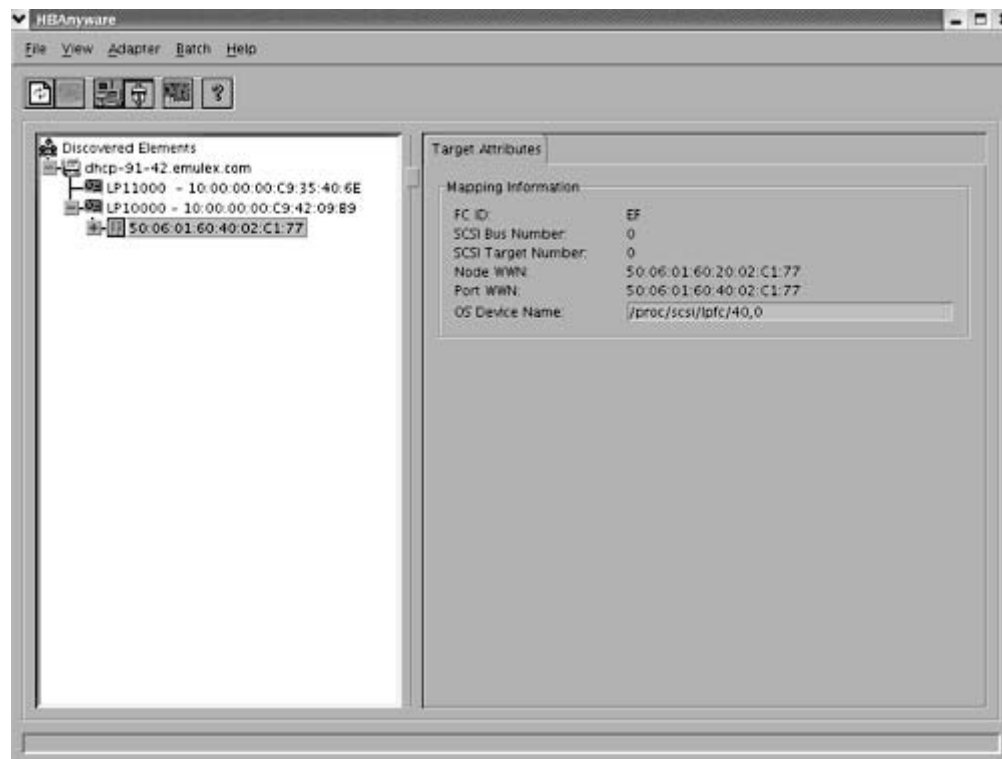


Figure 9: The HBAAnyware Target Attributes Tab


Target Attributes Field Definitions

- Mapping Information
 - FC ID - the Fibre Channel ID for the target; assigned automatically in the firmware.
 - SCSI Bus Number - defines the SCSI bus to which the target is mapped.
 - SCSI Target Number - the target's identifier on the SCSI bus.
 - Node WWN - the unique 64-bit number, in hexadecimal.
 - Port WWN - the unique 64-bit number, in hexadecimal.
 - OS Device Name - operating system device name.

Viewing LUN Attributes

The **LUN Attributes** tab in HBAAnyware contains information specific to the selected logical unit number (LUN).

To view the LUN attributes:

1. Start HBAAnyware.
2. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Host Name**.
 - From the toolbar, click the **Group HBAs by Host Name** button. 

3. Click a LUN in the discovery tree.

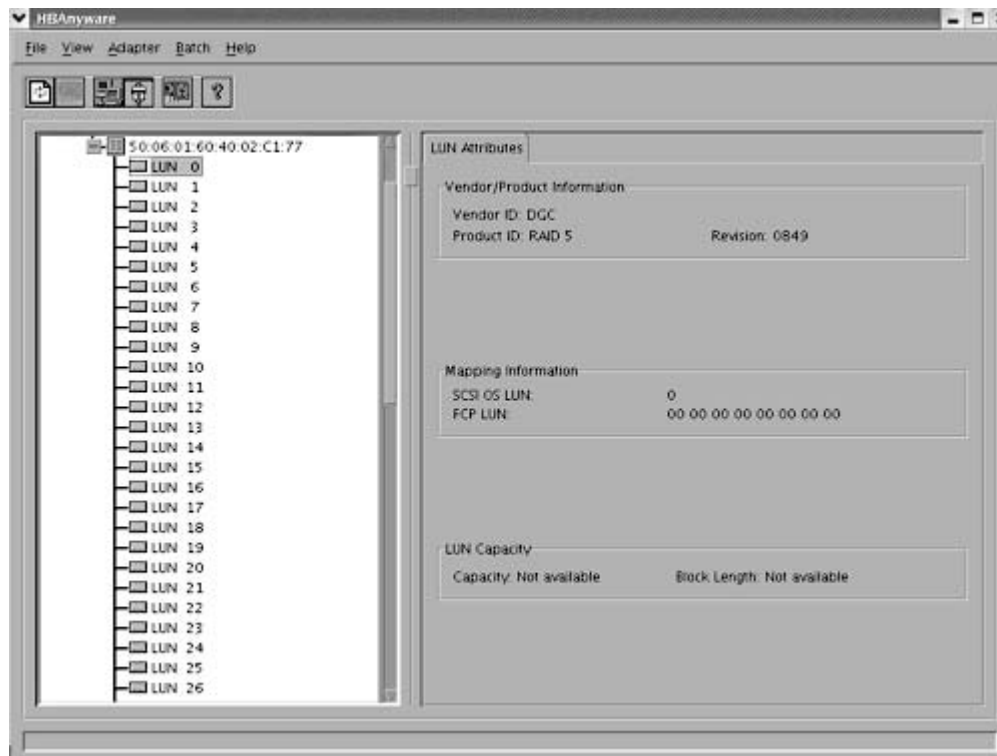


Figure 10: The HBAnyware LUN Attributes Tab

Lun Attributes Field Definitions

- Vendor Product Information
 - Vendor ID - the name of the vendor of the logical unit.
 - Product ID - the vendor-specific ID for the logical unit.
 - Revision - the vendor-specific revision number for the logical unit.

- Mapping Information
 - SCSI OS LUN - the SCSI identifier used by the operating system to map to the specific LUN.
 - FCP LUN - the Fibre Channel identifier used by the host bus adapter (HBA) to map to the SCSI OS LUN.
- LUN Capacity
 - Capacity - the capacity of the logical unit, in megabytes.
 - Block Length - the length of a logical unit block in bytes.

Viewing Fabric Attributes

The **Fabric Attributes** tab in HBAnyware contains information specific to the selected fabric.

To view the fabric attributes:

1. Start HBAnyware.
2. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Fabric Address**.
 - From the toolbar, click the  button.
3. Click on a fabric address in the discovery tree.

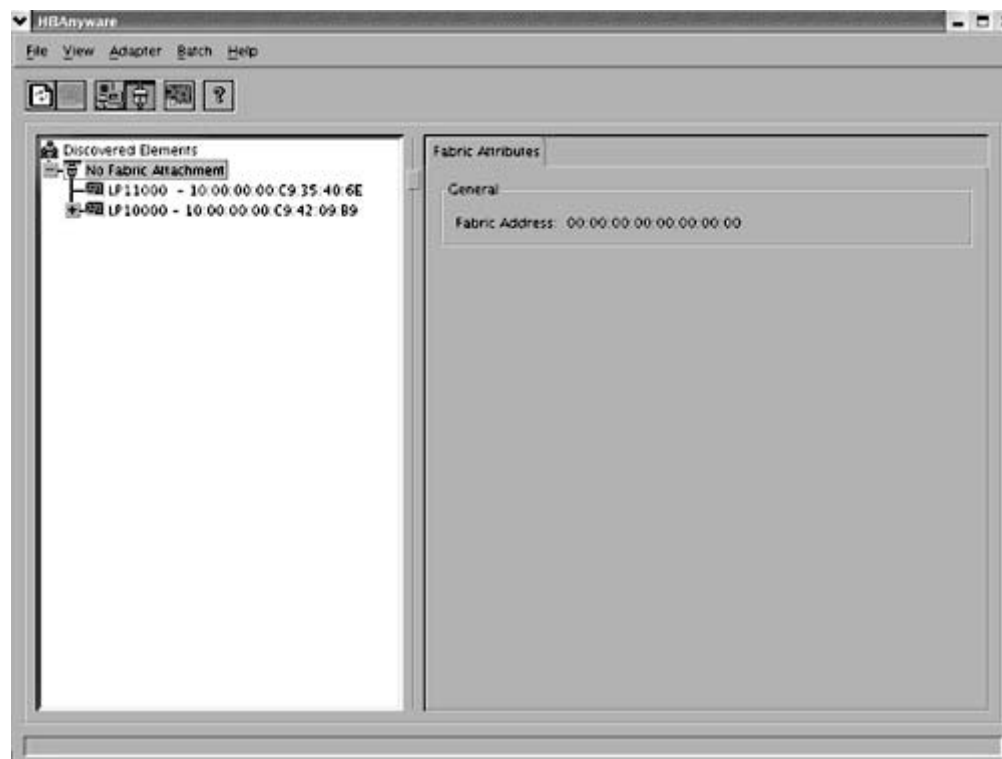


Figure 11: The HBAnyware Fabric Attributes Tab

General Area Field Definitions

- Fabric Address - a 64-bit unique identifier assigned to each Fibre Channel fabric.

Viewing General HBA Attributes

The **General** tab in HBAnyware contains general attributes associated with the selected host bus adapter (HBA).

To view general attributes:

1. Start HBAnyware.
2. If desired, sort the discovered HBAs.
3. Click an HBA in the discovery tree.

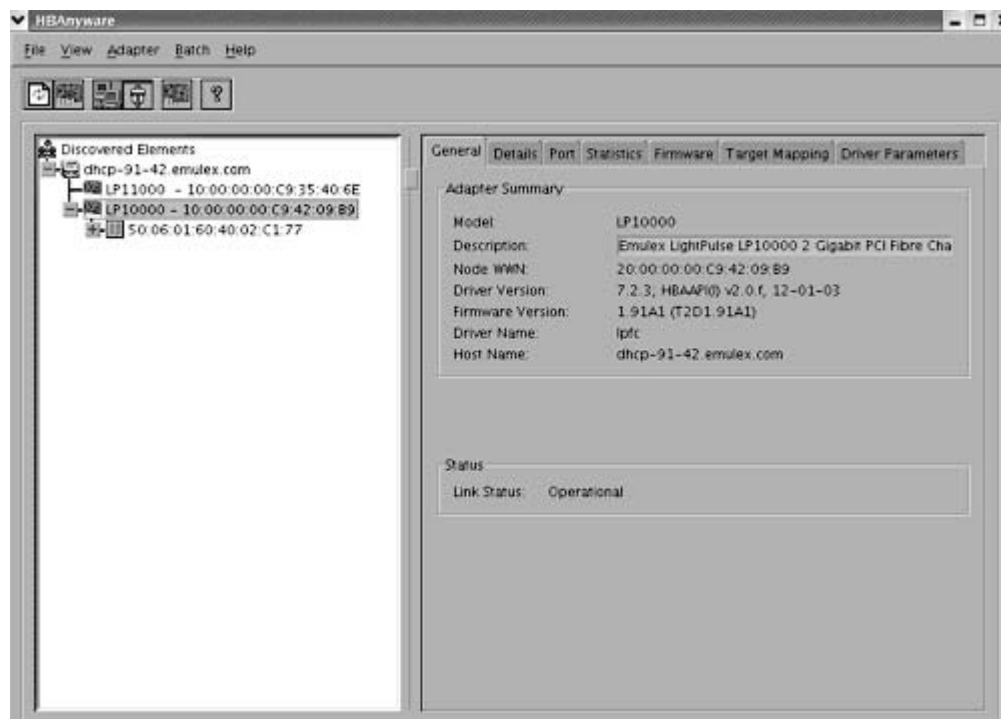


Figure 12: The HBAnyware General Attributes Tab

Adapter Summary Field Definitions

- Model - the Emulex HBA model number.
- Description - a formal description of the HBA, including model number, bus type and link speed. This field is recessed, indicating that the information in this field may exceed the visible length of the field. Use the arrow keys on your keyboard to scroll and view additional information.
- Node WWN - a 64-bit worldwide unique identifier assigned to the node.
- Driver Version - the driver version number and the HBA application programming interface (HBA API) version number.
- Firmware Version - the version of Emulex firmware currently active on the HBA.
- Driver Name - the executable file image name for the driver as it appears in the Emulex driver download package.

- Host Name - the name of the host to which the driver was downloaded.

Status Area Field Definitions

This field reflects the current state of the HBA. There are several possible states:

- The operational state indicates that the HBA is connected to the network and operating normally.
- All other states indicate that the HBA is not connected to the network. Green HBA icons with red descriptive text indicate that the HBA is offline. These offline states are:
 - User offline - the adapter is down or not connected to the network.
 - Bypassed - the HBA is in Fibre Channel discovery mode.
 - Diagnostic Mode - the HBA is controlled by a diagnostic program.
 - Link Down - there is no access to the network.
 - Port Error - the HBA is in an unknown state; try resetting it.
 - Unknown - the HBA is offline for an unknown reason.
 - Resetting - the HBA is in the process of rebooting.
 - Downloading - a firmware or other image is being downloaded to the HBA.

Viewing Detailed HBA Attributes

The **Details** tab in HBAnyware contains detailed attributes associated with the selected HBA.

To view the detailed attributes:

1. Start HBAnyware.
2. If desired, sort the discovered HBAs.
3. Click an HBA in the discovery tree. The **General** tab is displayed.

4. Click the **Details** tab.

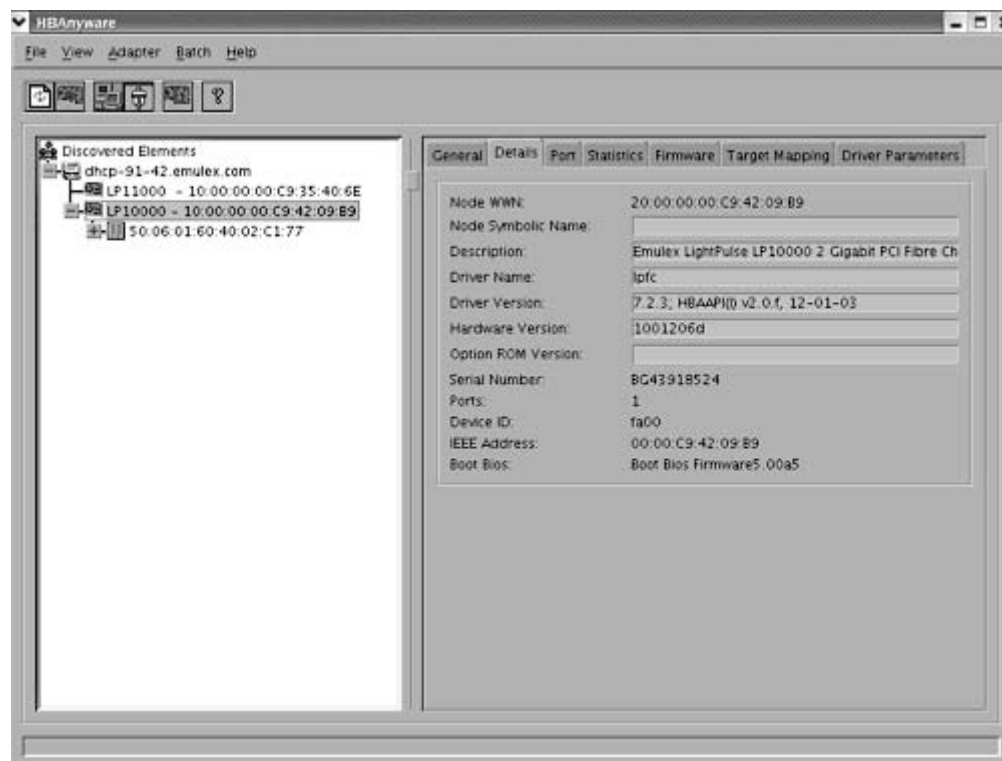


Figure 13: The HBAAnyware Details Tab

Note: Recessed fields indicate that the information in that field may exceed the text display area of the field. Use the arrow keys on your keyboard to scroll and view additional information.

Adapter Details Field Definitions

- Node WWN - a 64-bit worldwide unique identifier assigned to the node.
- Node Symbolic Name - in a fabric, the name registered with the name server.
- Description - a formal description of the HBA, including model number, bus type and link speed.
- Driver Name - an executable file image name for the driver as it appears in the Emulex driver download package.
- Driver Version - the driver version number and the HBA application programming interface (HBA API) version number.
- Hardware Version - the board version number, represented by the JEDEC ID, which is machine-readable from the Emulex Application Specific Intergrated Circuit (ASIC).
- Option ROM Version - the optional read-only memory version number; displayed if the BootBIOS bootup message is enabled on the HBA.
- Serial Number - the serial number assigned to the HBA when it was manufactured. Typically, this is a Binary Coded Decimal (BCD) string of the 48-bit IEEE address for the HBA.
- Ports - the number of ports on the HBA. Currently, this is always one. The two ports of dual-channel HBAs are displayed in the discovery tree as two HBAs.
- Device ID - the HBA's default device ID.

- IEEE Address - the Media Access Control (MAC) address is in conformance with the Fibre Channel Link Encapsulation (FC-LE) standard. This address is a 48-bit number that is unique to every HBA in existence. The IEEE Address is printed on a label affixed to one end of the HBA.

Viewing Port Information

The **Port** tab in HBAnyware contains information about the port on the selected host bus adapter (HBA).

To view port information:

1. Start HBAnyware.
2. If desired, sort the discovered HBAs.
3. Click an adapter in the discovery tree. The **General** tab is displayed.
4. Click the **Port** tab.

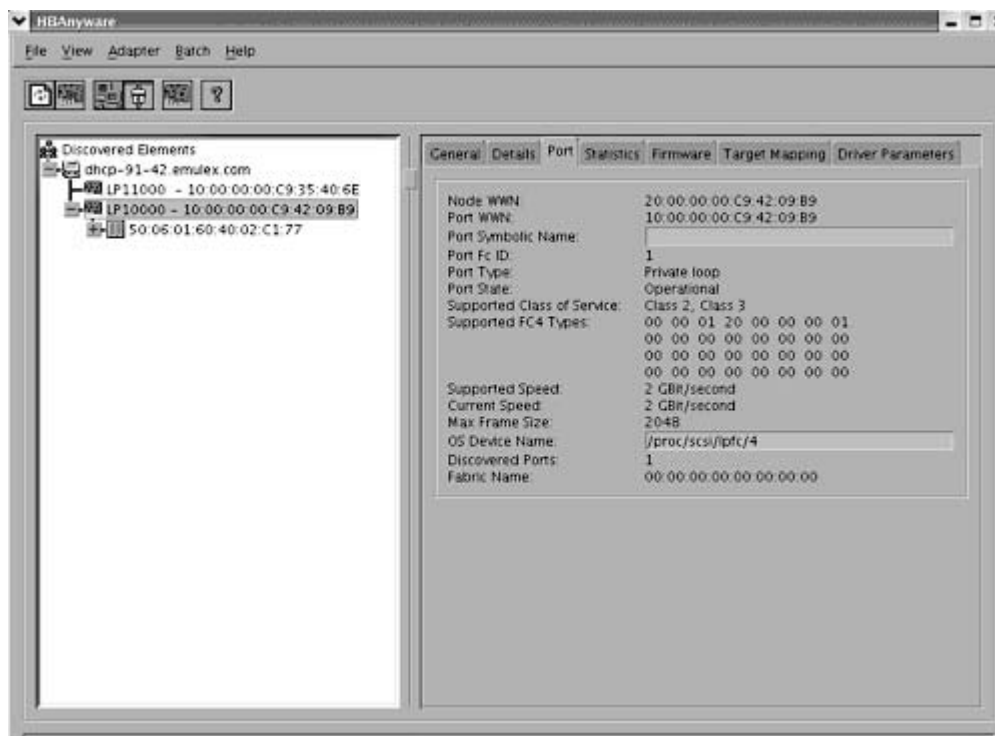


Figure 14: The HBAnyware Port Tab

Port Field Definitions

- Node WWN - a 64-bit worldwide unique identifier assigned to the node. The Node WWN is communicated during the login and port discovery processes. This identifier stays with the entity for its lifetime.
- Port WWN - a 64-bit worldwide unique identifier assigned to the port. The Port WWN is communicated during the login and port discovery processes. This identifier stays with the entity for its lifetime.
- Port Symbolic Name - the name registered by the HBA with a name server. This field is recessed, indicating that the information in this field may exceed the visible length of the field. Use the arrow keys on your keyboard to scroll and view additional information.
- Port Fc ID - Fibre Channel ID for the port.

- Port Type - describes the current operational mode of the port.
- Port State - current status of the port: operational or link down.
- Supported Class of Service - a frame delivery scheme exhibiting a set of delivery characteristics and attributes. There are three classes of service.
 - Class-2 provides a frame switched service with confirmed delivery or notification of non-delivery.
 - Class-3 provides a frame switched service similar to Class-2 but without notification of frame delivery or non-delivery.
- Supported FC4 Types - a 256-bit (8-word) map of the FC-4 protocol types supported by the port. Each bit in the map corresponds to a Type value as defined by the Fibre Channel standards and contained in the Type field of the frame header.
- Supported Speed - maximum link speed supported by the HBA.
- Current Speed - link speed for the current session.
- Max Frame Size - maximum frame size.
- OS Device Name - the platform-specific name by which the HBA is known to the operating system.
- Discovered Ports - number of facilities that provide Fibre Channel interface attachment.
- Fabric Name - 64-bit worldwide unique identifier assigned to the fabric.

Viewing Statistics

The **Statistics** tab in HBAnyware provides cumulative totals for various error events and statistics on the port. Statistics are cleared when the host bus adapter (HBA) is reset.

To view statistics:

1. Start HBAnyware.
2. If desired, sort the discovered HBAs.
3. Click an HBA in the discovery tree. The **General** tab is displayed.

- Click the **Statistics** tab.

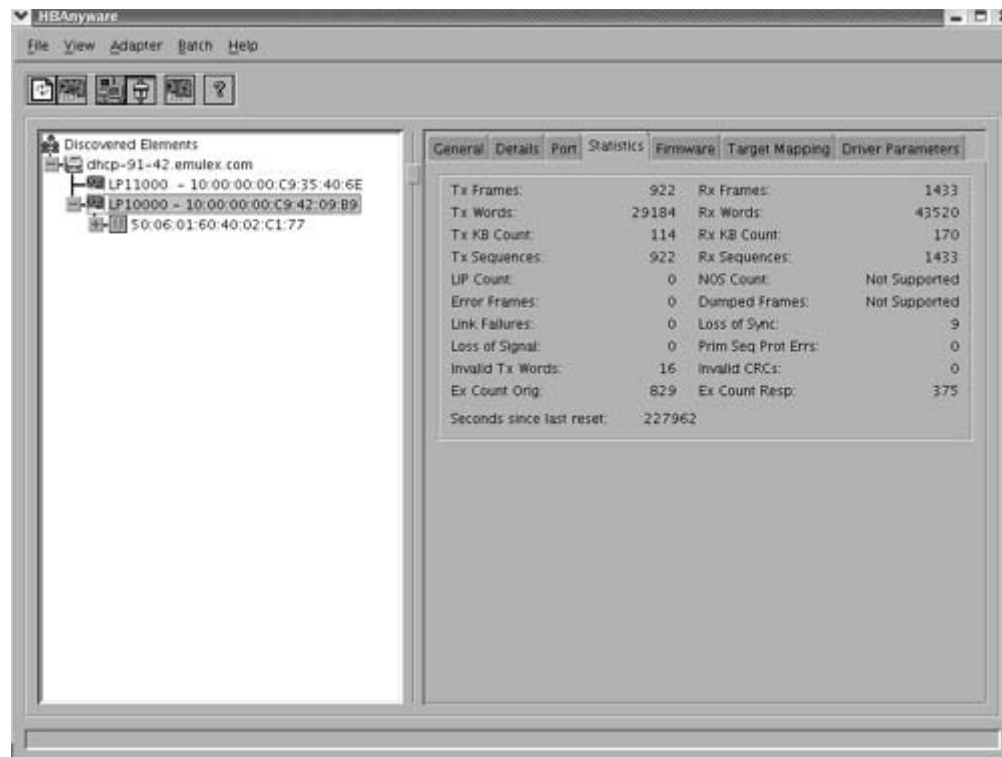


Figure 15: The HBAnyware Statistics Tab

Statistics Field Definitions

- Tx Frames - Fibre Channel frames transmitted by this HBA port.
- Tx Words - Fibre Channel words transmitted by this HBA port.
- Tx KB Count - Fibre Channel kilobytes transmitted by this HBA port.
- Tx Sequences - Fibre Channel sequences transmitted by this HBA port.
- Link Failures - the number of times the link failed. A link failure is a possible cause of a timeout.
- Loss of Signal - the number of times the signal was lost.
- Invalid Tx Words - the total number of invalid words transmitted by this HBA port.
- Ex Count Orig - the number of Fibre Channel exchanges originating on this port.
- LIP count - the number of loop initialization primitive (LIP) events that have occurred for the port. This field is not supported if the topology is not arbitrated loop. Loop initialization consists of the following:
 - Temporarily suspend loop operations.
 - Determine whether loop capable ports are connected to the loop.
 - Assign AL_PA IDs.
 - Provide notification of configuration changes and loop failures.
 - Place loop ports in the "monitoring" state.
- NOS count - this statistic is currently not supported.
- Rx Frames - the number of Fibre Channel frames received by this HBA port.

- Rx Words - the number of Fibre Channel words received by this HBA port.
- Rx KB Count - the received kilobyte count by this HBA port.
- Rx Sequences - the number of Fibre Channel sequences received by this HBA port.
- Loss of Sync - the number of times loss of synchronization has occurred.
- Prim Seq Prot Errs - the primitive sequence protocol error count. This counter is incremented whenever there is any type of protocol error.
- Invalid CRCs - the number of frames received that contain CRC failures.
- Ex Count Resp - the number of Fibre Channel exchange responses made by this port.
- Error Frames - the number of frames received with cyclic redundancy check (CRC) errors.
- Dumped Frames - this statistic is not currently supported.
- Seconds Since Last Reset - the number of seconds since the HBA was last reset.

Viewing Firmware Information

Use the **Firmware** tab to view current firmware versions and update firmware on remote and local HBAs. The update procedure is on page 54.

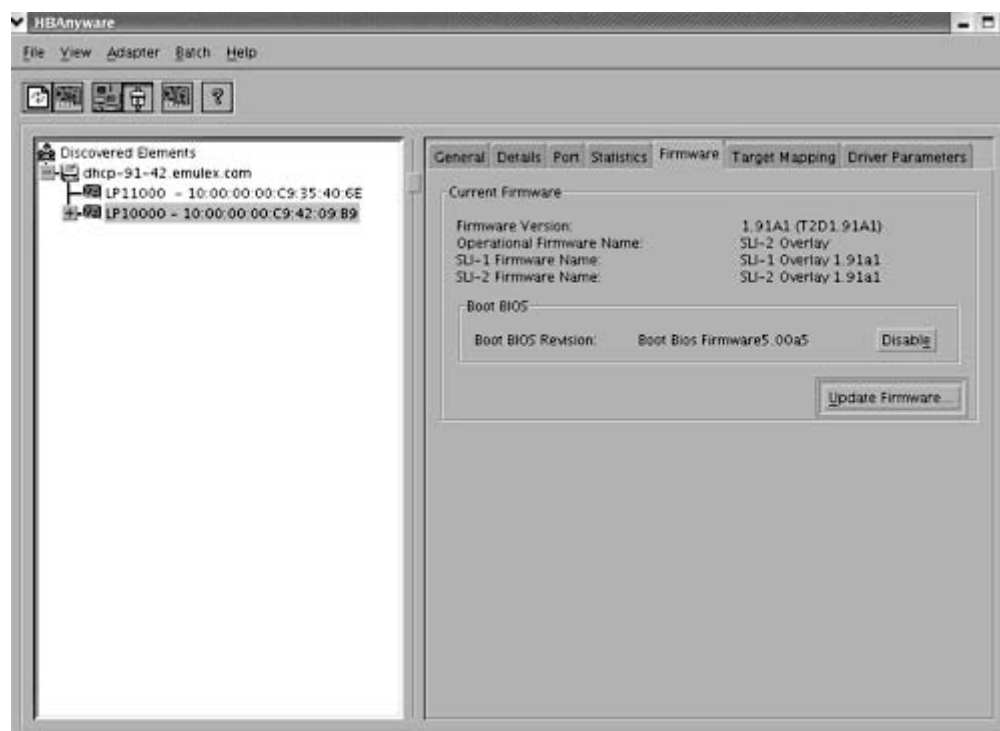


Figure 16: The HBAAnyware Firmware Tab

Firmware Field Definitions

- Firmware Version - the Emulex firmware version number for this model of HBA.
- Operational Firmware Name - if visible, the name of the firmware that is operational.
- SLI-1 Firmware Name - the name of the SLI-1 firmware overlay.
- SLI-2 Firmware Name - the name of the SLI-2 firmware overlay.

- Boot BIOS - when present, displays revision of the Boot BIOS firmware.

Firmware Tab Buttons

- **Enable/Disable** - click to enable or disable the Boot BIOS for the HBA.

Note: If the state of the boot code message on the board has changed, this change will be reflected immediately on the Details tab.

- **Update Firmware** - click to this button to display the **HBAnyware Firmware Download** dialog box. Using the **HBAnyware Firmware Download** dialog box, you can browse to the firmware file you wish to download and download the file. Refer to the “Update Firmware Using HBAnyware” topic on page 54 for more information.

Viewing Target Mapping

Use this tab to perform mapping and persistent binding tasks. See page 75 to learn how to set up persistent binding.

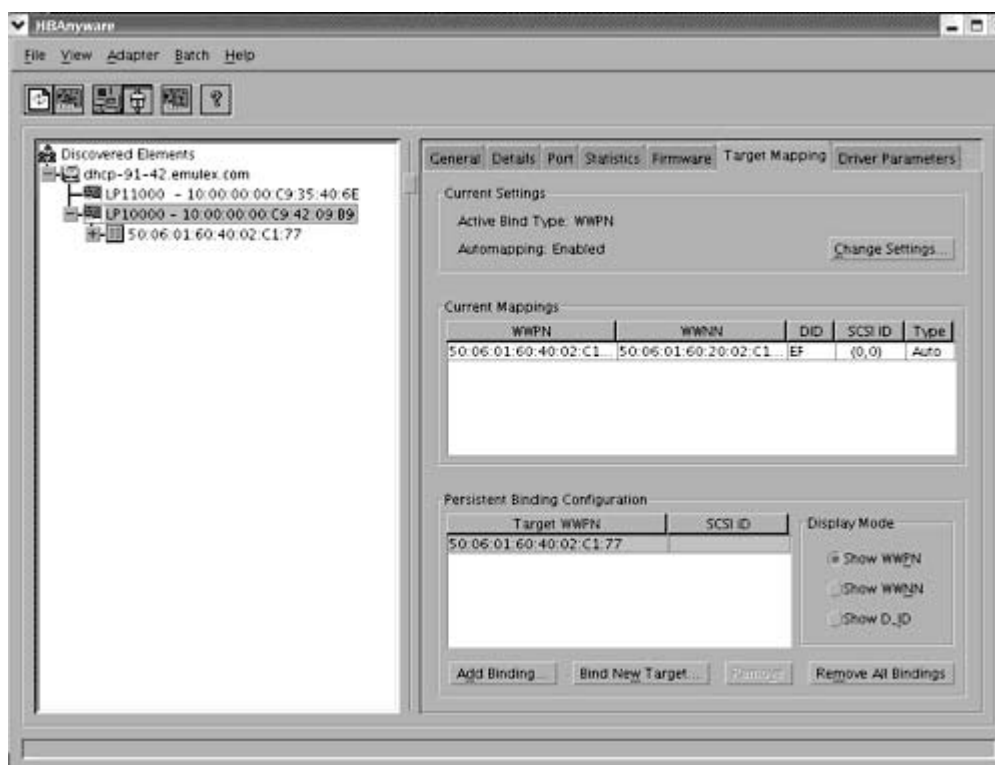


Figure 17: The HBAnyware Target Mapping Tab

Target Mapping Field Definitions

Current Settings Area

- Active Bind Type - displays the currently active bind type: world wide port name (WWPN), world wide node name (WWNN), or a destination identifier (D_ID).
- Automapping - current state of SCSI device automapping: enabled (default) or disabled.

Current Mappings Table

- This table lists all currently mapped targets for the selected HBA.

Persistent Binding Configuration Table

- This table lists persistent binding information for the selected HBA.

Display Mode Radio Buttons

- Show WWPN
- Show WWNN
- Show D_ID


Target Mapping Buttons

- **Change Settings** - click to change the Bind Type, the mode used to persistently bind target mappings. The **Mapped Target Settings** window is displayed. Select the Bind Type (WWPN, WWNN, or D_ID) or set Automapping to Enabled to Disabled.
- **Add Binding** - click to add a persistent binding.
- **Bind New Target** - click to add a target that does not appear in the Persistent Binding table.
- **Remove** - click to remove the selected binding.
- **Remove All Bindings** - click to remove all persistent bindings that are currently defined for the selected HBA.

Viewing Driver Parameters

The **Driver Parameters** tab allows you to view and modify driver parameters either for an individual HBA or for all adapters, with the same single driver type and version, that are in one host.

To display the driver parameters for an adapter:

1. Start HBAnyware.
2. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Host Name**.
 - From the toolbar, click the **Group HBAs by Host Name** button. 
3. In the discovery tree, click the adapter for which you want to change a parameter. The **General** tab is displayed.
4. Click the **Driver Parameters** tab. The Installed Driver Types field displays the driver operating system version that is installed on the HBA (Figure 18).

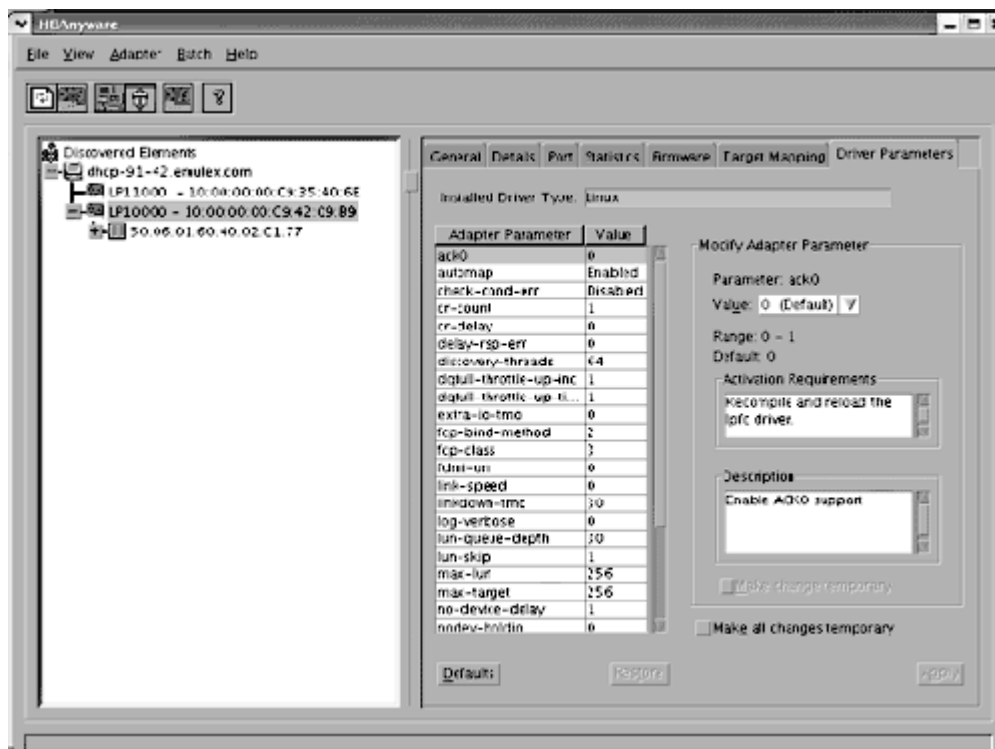


Figure 18: The HBAnyware, HBA Selected, Driver Parameters Tab

To display the driver parameters for a host:

1. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Host Name**.
 - From the toolbar, click the  button.
2. In the discovery tree, click the host for which you want to change a parameter. The **Host Attributes** tab is displayed.

3. Click the **Host Driver Parameters** tab (Figure 19). The Installed Driver Types drop-down box displays a list of all driver types and driver versions that are installed on the adapters in the host.

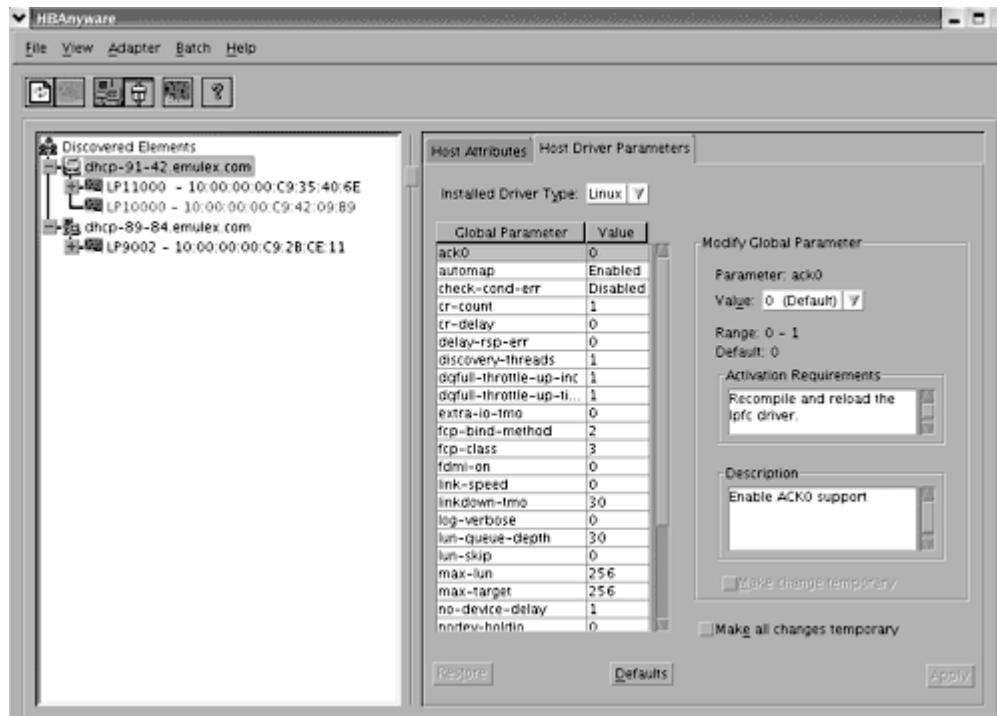


Figure 19: The HBAAnyware, Host Driver Parameters Tab

Driver Parameter Tab Field Definitions

- Installed Driver Type - current driver and version installed.
- Adapter Parameter table - a list of parameters and their current values.
- Parameter-specific information - details about the parameter appears on the right side of the tab.

Driver Parameter Tab Buttons

- **Restore** - click to restore parameters to their original values, (i.e., the values they had before you made changes).
- **Defaults** - click to reset all parameter values to their default (out-of-box) values.
- **Apply** - click to apply any driver parameter changes. If you changed a parameter that is not dynamic, you must reboot.

Note: For a procedure on changing a driver parameter's value using HBAAnyware, see page 59.

Viewing HBA Information using lputil

The LightPulse Diagnostic utility (lputil) allows you to view information for a selected adapter.

To view HBA information using lputil:

1. Start lputil, the Main menu opens:

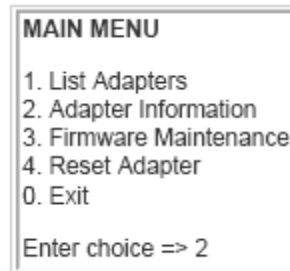


Figure 20: The Main Menu

2. Select choice #2.

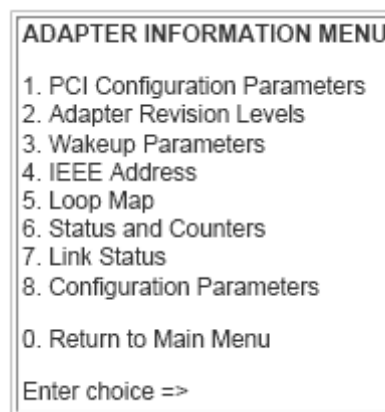


Figure 21: The Adapter Menu

If you have multiple adapters, a list is displayed, you select an adapter and the Adapter Information Menu opens. If you have only one adapter, the Adapter Information Menu opens for that adapter.

- PCI Configuration Parameters - parameters from the PCI configuration space on the adapter. Examples of this information include vendor ID, device ID, base addresses, ROM address, header type, subclass and base class.
- Adapter Revision Levels - firmware revision levels, including kernel and overlay version information.
- Wakeup Parameters - BIOS status and version, as well as SLI (service level interface).
- IEEE Address - the adapter board address.
- Loop Map - If you are currently using arbitrated loop topology, this menu option displays information about your connected devices, such as AL_PA and D_ID.
- Status and Counters - byte, frame, sequence and busy counts.
- Link Status - tracks activities such as link failure, loss of sync, and elastic overlay.
- Configuration Parameters - D_ID topology, and timeout values for link failures and loss of sync.

Resetting Adapters


You can reset adapters using either HBAnyware or lputil.

- HBAnyware allows you to reset remote and local adapters.
- lputil allows you to reset local adapters only.

Resetting the HBA Using HBAnyware

Caution: Do not reset your adapter while copying or writing files. This could result in data loss or corruption.

To reset a host bus adapter:

1. Start HBAnyware.
2. In the directory tree, click the HBA you want to reset.
3. Do one of the following:
 - From the menu bar, click **Adapter**, and then click **Reset Adapter**.
 - Click the **Reset** toolbar button. 
4. A warning window appears asking the user if they want to continue. Click **Yes**.

Resetting the HBA Using lputil

Caution: Do not reset your adapter while copying or writing files. This could result in data loss or corruption.

The LightPulse utility (lputil) allows you to reset the adapter.

To reset the adapter using lputil:

1. Start lputil. The Main menu is displayed (Figure 20).
2. Choose #4 Reset Adapter.
3. If you have multiple adapters, select the adapter you want to reset.

Resetting the adapter runs self tests and reestablishes links (causes discovery of devices). Once the adapter has been successfully reset, the Main menu is displayed.

Updating Firmware

You can update firmware using either HBAnyware or lputil.

- HBAnyware allows you to update firmware on remote and local HBAs.
- lputil allows you to update firmware on local HBAs only.

Updating Firmware Using HBAnyware

Prerequisites

- The firmware file has been downloaded from the Emulex Web site to the Emulex Repository folder (RMRepository). This folder is in /usr/sbin/hbanyware/RMRepository.
- The firmware file has been extracted into the Emulex Repository folder. When updating firmware on a remote system, the firmware is automatically transferred to the remote system and placed in the Emulex Repository folder (RMRepository).

Procedure

To load firmware using HBAnyware:

1. Start HBAnyware.
2. In the discovery tree (left pane), click the adapter to which you want to load the firmware.
3. In the property tabs (right pane), select the **Firmware** tab.

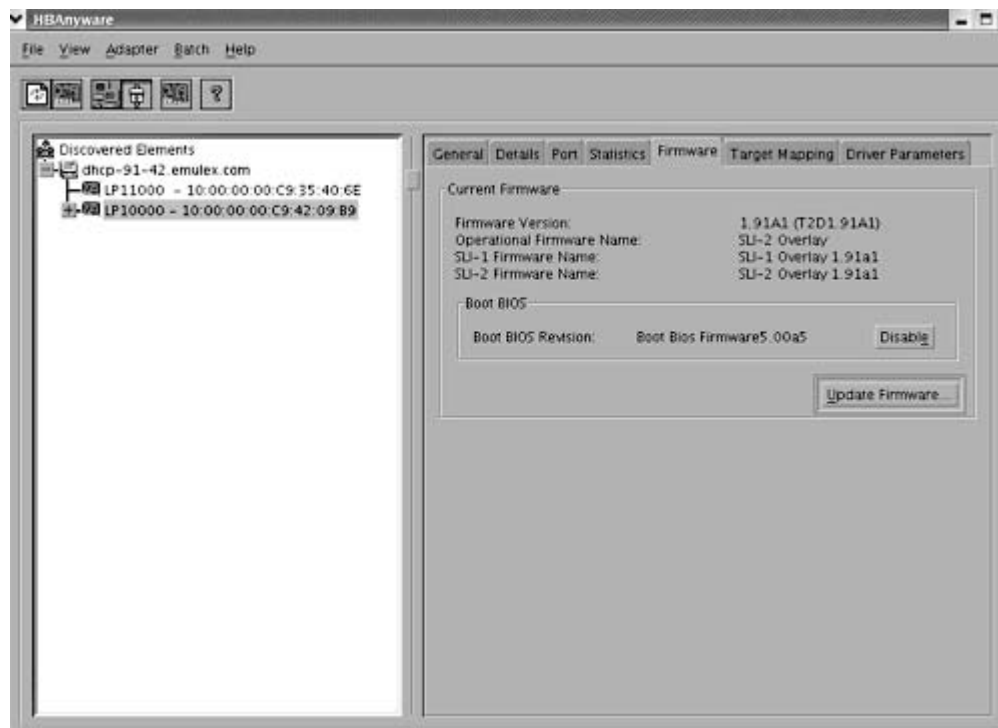


Figure 22: The HBAnyware Firmware Tab

- On the **Firmware** tab, click the **Update Firmware** button. The **HBAnyware Firmware Download** dialog box is displayed.




Figure 23: The HBAnyware Firmware Download dialog box

- Click **Browse**. The **Firmware File Selection** dialog box appears. Select the firmware file to download and click **OK**. The **HBAnyware Firmware Download** dialog box reappears.
- Click **Start Download**. A status bar displays the progress of the download. During this time the host bus adapter (or adapters) in the discovery tree is displayed in red text, indicating that it is offline. It is displayed in black text when the update is complete.

If you are updating the firmware on a dual-channel HBA, repeat steps 2 through 5 to update the firmware on the second port or use the “Batch Firmware Download” procedure on page 55.

Note: If the state of the boot code message on the board has changed, this change will be reflected immediately on the **Details** tab.

Loading Firmware (Batch Mode) Using HBAnyware

Downloading firmware in batch mode allows you to install firmware on multiple HBAs in a single step. Batch firmware loading is restricted to a single firmware file.

Note: No other HBAnyware functions can be performed while batch firmware loading is in progress.

Prerequisites

- The firmware file has been downloaded from the Emulex Web site and extracted to the Emulex Repository folder (RMRepository). This folder is in /usr/sbin/HBAnyware/RMRepository.

Procedure

To load firmware using batch mode:

- From the menu bar, select **Batch** and click **Download**.

Note: You do not need to select a particular tree element for this operation.

- The **Batch Firmware Download** dialog box is displayed.

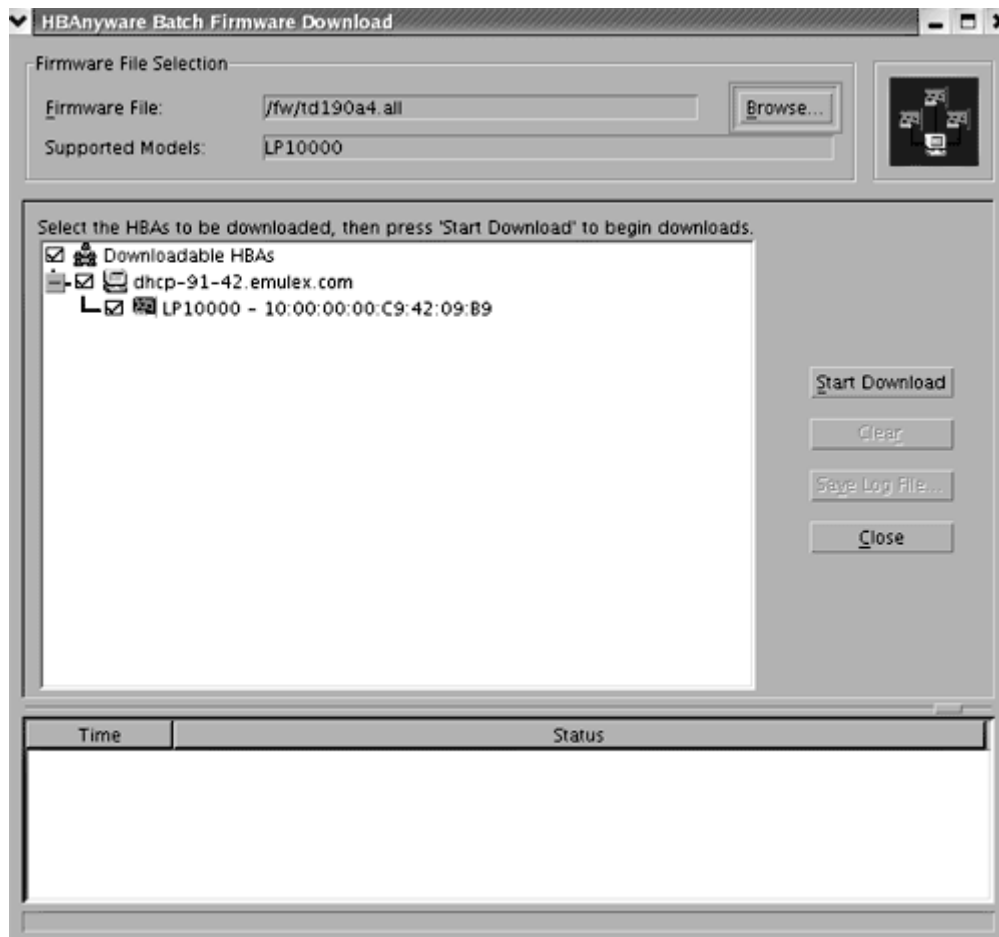


Figure 24: The HBAAnyware Batch Firmware Download dialog box

- Click **Browse**. The **Firmware File Selection** dialog box appears. Select the firmware file to download and click **OK**. The **HBAAnyware Batch Firmware Download** dialog box reappears. A tree-view appears showing all HBAs and their corresponding hosts for which the selected firmware file is compatible.
- Click the check box next to an HBA or host to select or remove that HBA from the batch process. (Figure 24)
- Click **Start Download**. The Status pane displays the progress of the download. After downloading begins, the tree-view displays the progress. As a file for a selected HBA is downloading, it appears orange in the tree-view. After completion, the entry for the HBA changes to green for a successful download or red if it failed.
- When downloading is complete, click **Save Log File** if you wish to view or print the log file.
- Click **Close** to exit the batch procedure.

Updating Firmware Using lputil

Prerequisites

- The driver for Solaris (including lputil) is installed properly.
- The firmware file has been downloaded to a local drive.

This procedure uses the lputil utility, which is installed with the driver.

Procedure

Caution: Do not interrupt this process or power down the system until the process is complete.

To update firmware using lputil:

1. Start the utility by entering the complete path to lputil. The path in the example reflects the default installation path. If the installation path was modified, adjust the command appropriately.
`/usr/sbin/lpfc/lputil`
2. From the Main menu, enter 3, Firmware Maintenance.
3. If prompted, choose the HBA that is being updated.
4. Enter 1, Load Firmware Image.
5. Enter the full path to the firmware file.
6. Enter 0 twice to exit the utility.

The new firmware is transferred to flash ROM.

If you are updating the firmware on a dual-channel HBA, repeat steps 3 through 5 to update the firmware on the second port.

Updating FC Boot

You can update FC boot code (BootBIOS, OpenBoot or EFIBoot) using either HBAnyware or lputil.

- HBAnyware allows you to update FC boot code on remote and local HBAs.
- lputil allows you to update FC boot code on local HBAs only.

Updating FC Boot Code Using HBAnyware

Prerequisites

- The Emulex driver for Solaris has been installed properly.
- HBAnyware has been installed properly.
- The FC Boot files have been downloaded from the Emulex Web site to the Emulex Repository folder (RMRepository). This folder is in `/usr/sbin/hbanyware/RMRepository`.
- The FC Boot file has been extracted into the Emulex Repository folder.

Caution: If you are downloading EFIBoot on an HBA attached to the remote system disk, it is recommended to use the EFI utility to perform the download. The EFI utility is bundled with the EFI boot firmware on the Emulex website.

Procedure

To update FC Boot Code using HBAware:

1. Start HBAware.
2. In the discovery tree (left pane), click the adapter to which you want to load the firmware.
3. In the property tabs (right pane), select the **Firmware** tab.
4. On the **Firmware** tab, click the **Update Firmware** button. The **Firmware Download** window is displayed.
5. Browse to the Emulex Repository. Select the EFIBoot file to download and click **OK**. Click the **Start Download** button. A status bar displays the progress of the download. During this time the host bus adapter (or adapters) in the discovery tree is displayed in red text, indicating that it is offline. It is displayed in black text when the update is complete.
6. Reboot the system.

If you are updating EFIBoot on a dual-channel HBA, repeat steps 2 through 5 to update the EFI-Boot on the second port or refer to the “Batch Firmware Download” procedure on page 55.

Note: If the state of the boot code message on the board has changed, this change will be reflected immediately on the **Details** tab.

Updating FC Boot Code Using lputil

Prerequisites

- The driver for Solaris is installed properly.
- The boot code file has been downloaded to a local drive.

Procedure

To update FC Boot Code using lputil:

1. Start the utility by entering the complete path to lputil. The path in the example reflects the default installation path. If the installation path was modified, adjust the command appropriately.
 - Enter the following command:

```
/usr/sbin/lpfc/lputil
```
2. From the Main menu, enter 3, Firmware Maintenance.
3. Enter 1, Load Firmware Image.
4. Enter the full path to the boot code file.

The new boot code is transferred to flash ROM.

5. Enter 0 twice to exit.

Enabling/Disabling BootBIOS Using lputil

Prerequisites

- The Emulex driver for Solaris is installed properly.
- To enable BootBIOS, the BIOS file has been downloaded to a local drive.

Procedure

To enable or disable BootBIOS using lputil:

1. Start the utility by entering the complete path to lputil. The path in the example reflects the default installation path. If the installation path was modified, adjust the command appropriately.

```
/usr/sbin/lpfc/lputil
```

2. From the Main menu, enter 3, Firmware Maintenance.

The Firmware Maintenance menu is displayed.

3. From the Firmware Maintenance menu, press 6, Boot BIOS Maintenance.

- If the boot code is currently disabled, press 1, Enable Boot BIOS, to enable the boot code.
- If the boot code is already enabled, press 1, Disable Boot BIOS, to disable the boot code.
- If the boot code is not currently loaded, the following message is displayed:

```
There is no Boot BIOS found on adapter
```

4. Enter 0 twice to exit.

Configuring the Driver

You can configure the driver using HBAnyware or lpfc.conf.

Note: HBAnyware enables you to make dynamic parameter changes with any version of Solaris. However, changes made by editing the lpfc.conf file and issuing an update_drv command are only dynamic for Solaris 9 and later.

The **Driver Parameters** tab allows you to modify driver parameters either for an individual adapter or for all adapters, with the same single driver type and version, that are in one host. For example, if you set driver parameters on a host that includes two adapters, you can make changes to the driver parameters for both adapters simultaneously using the **Host Driver Parameters** tab.

For each parameter, the tab displays the current value, the range of acceptable values, the default value, and whether the parameter is dynamic (a dynamic parameter allows the change to take effect without restarting the HBA or rebooting the system).

Configuring Driver Parameters Using HBAnyware

Change a Parameter's Value

To change a parameter's value using HBAnyware:

1. View the driver parameters for the HBA or the host.

- In the appropriate parameter tab, click the parameter that you want to change. Information about the parameter appears on the right side of the screen.

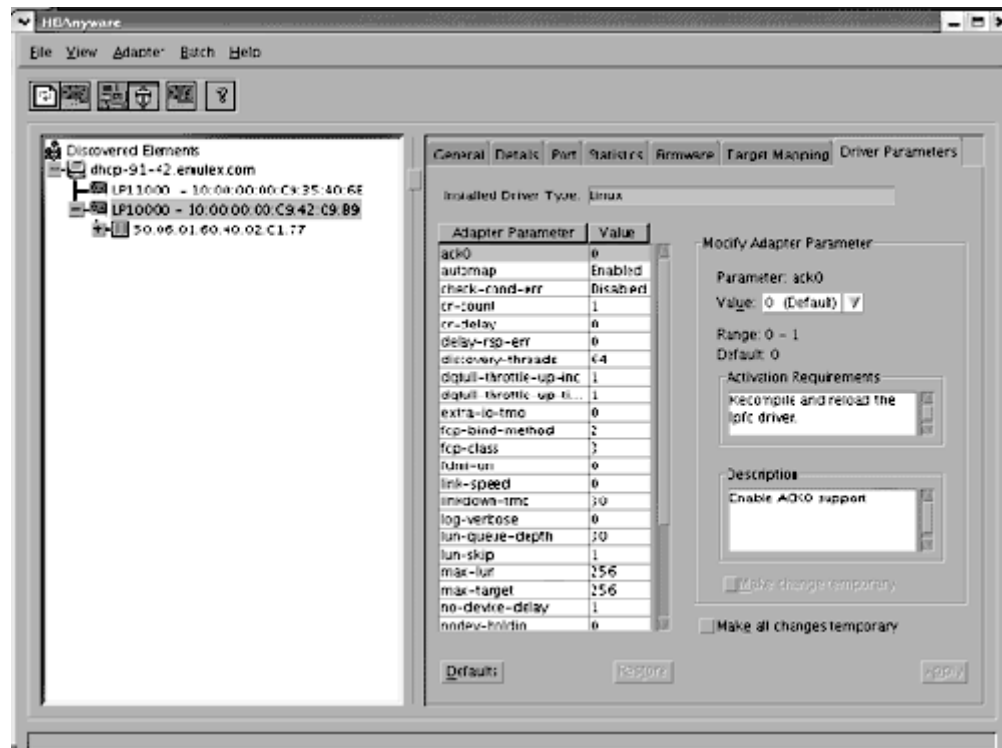


Figure 25: The HBAnyware, HBA Selected, Driver Parameters Tab

- Change to the new value. Some parameters contain radio buttons to enable or disable a parameter, others include pull down values and some contain a blank field in which you enter the new value.
- If you want the change to be temporary (causing the parameter to revert to its last permanent setting when the system is rebooted), check the **Make change temporary** box. This option is available only for dynamic parameters.
- If you are making changes to multiple parameters, and you want all the changes to be temporary, check the **Make all changes temporary** box. This setting overrides the setting of the "Make change temporary" box. Only dynamic parameters can be made temporary.
- To apply your changes, click **Apply**.
- For parameters which cannot be changed dynamically (link speed for example), the system must be rebooted.

Restoring All Parameters to Their Earlier Values

If you have made changes to parameters and have not saved them by clicking **Apply**, and you want to restore the parameters to their last saved values, click **Restore**.

Resetting All Default Values

Click **Defaults** to reset all parameter values to their default (out-of-box) values.

Configuring Driver Parameters Using lpfc.conf

Using lpfc.conf for Solaris 9 and 10

Typically, the lpfc driver reads the configuration parameters from lpfc.conf, during driver startup. Once the configuration parameters are read, the lpfc driver does not re-read the driver parameters. This feature is native to the Solaris operating system. However on Solaris releases 9 or later, you can force the lpfc driver to re-read the lpfc.conf configuration file using the update_drv system command.

Note: The update_drv command does not change the parameters for the lpfc instances that are already active in the system. The changes will be applicable to any new lpfc instances that are attached afterwards. Changes made by editing the lpfc.conf file and issuing an update_drv command are only dynamic for Solaris 9 and later.

To modify the configuration parameters:

1. Login as root or su to root.
2. Edit the configuration parameters in the file /kernel/drv/lpfc.conf. To view a table of all lpfc.conf parameters, see “The Configuration File (lpfc.conf)” on page 62.
3. Run update_drv -f lpfc to force the Solaris system to re-read lpfc.conf.

To modify parameters for an adapter instance that is already active:

1. Login as root or su to root.
2. Run cfgadm and note down the <Ap_Id> corresponding to the lpfc instance.
3. Unconfigure the lpfc driver instance corresponding to <Ap_Id>.
4. Edit the adapter-specific parameters in the file /kernel/drv/lpfc.conf.
5. Run update_drv -f lpfc to force the Solaris system to re-read lpfc.conf.
6. Run cfgadm and configure the adapter instance back.

To modify parameters for an adapter that is yet to be added:

1. Login as root or su to root.
2. Run cfgadm and configure the new adapter instance. Write down the adapter instance. The adapter instance is printed in the driver sign-on messages.
3. Run cfgadm and unconfigure the new adapter instance.
4. Edit the adapter-specific parameters in the file /kernel/drv/lpfc.conf.
5. Run update_drv -f lpfc to force the Solaris system to re-read lpfc.conf.
6. Run cfgadm and configure the new adapter instance.

Updating Parameters for Solaris 2.6, 7 and 8

To change parameters or bindings in Solaris 2.6, 7 or 8:

1. Quiesce all I/O on the device.
2. Unconfigure all ports with open instances to the driver.
3. Edit the adapter-specific parameters in the file /kernel/drv/lpfc.conf.
4. Unload the driver using the modunload command. (See “Loading or Unloading the Driver Without Rebooting” on page 31.)

5. Reload the driver using the modload command. (See“Loading or Unloading the Driver Without Rebooting” on page 31.)

Understanding Device Numbering

When you configure the driver, there are various configuration parameters that rely upon device numbers. This section will explain the two different levels of device numbering and how they apply to specific configuration parameters.

The first level of device numbering is the driver instance number. This is the number that is displayed in log messages to syslog or the console, for example, lpfcX. The lpfc.conf configuration parameters that use lpfcXtY rely on X to be the driver instance number and Y to be the FCP target number. This format is used for configuring LUN throttles and persistent binding. In systems where adapters are moved around or changed, it is possible for the driver instance number to exceed the number of adapters.

The second level of numbering is the SCSI controller number. This number is assigned by the system for each SCSI target driver it detects. It is typically in the special files created to access each SCSI device, for example /dev/dsk/cXt0d0s2. Once these numbers are assigned, they will stay the same between reboots if there are no HBA changes to the system.

The Configuration File (lpfc.conf)

The lpfc.conf file contains all the parameters that control driver initialization.

In the lpfc.conf file, all adapter-specific parameters have lpfcX-prefix (where X is the driver instance number); e.g., setting lpfc0-lun-queue-depth= 20 makes 20 the default number of maximum commands which can be sent to a single logical unit (disk).

Changes to the lpfc.conf file require you to unload and reload the driver.

Note: If you want to override a driver parameter for a single driver-loading session, you can specify it as a parameter to the modload command.
 For example: # modload /kernel/drv/lpfc automap=0 (for 32-bit platforms) or modload /kernel/drv/sparcv9/lpfc automap=0 (for 64-bit platforms) . This will load Emulex's SCSI support driver with automap set to 0 for this session.

Table 2: lpfc.conf Parameters

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|-----------|---------------------|---------|-------|------|---------|--|
| ack0 | Controller Specific | 0 | 0=Off | 1=On | No | Use ACK0 for class 2. If ack0 is 1, the adapter will try to use ACK0 when running Class 2 traffic to a device. If the device doesn't support ACK0, then the adapter will use ACK1. If ack0 is 0, only ACK1 will be used when running Class2 traffic. |

Table 2: Ipfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|-------------------|---------------------|---------|-------|------|---------|--|
| automap | Controller Specific | 1 | 0=Off | 1=On | No | Automatically assign SCSI IDs to FCP targets detected. If automap is 1, SCSI IDs for all FCP nodes without persistent bindings will be automatically generated based on the bind method of the corresponding HBA port. If FCP devices are added to or removed from the Fibre Channel 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 automap is 0, only devices with persistent bindings will be recognized by the system. |
| cr-count | Controller Specific | 1 | 1 | 255 | No | This value specifies a count of I/O completions after which an interrupt response is generated. This feature is disabled if cr-delay is set to 0. |
| cr-delay | Controller Specific | 0 | 0 | 63 | No | This value specifies a count of milliseconds after which an interrupt response generated if cr-count has not been satisfied. This value is set to 0 to disable the Coalesce Response feature as default. |
| delay-rsp-err | Controller Specific | 0 | 0=Off | 1=On | Yes | (Boolean) The driver will delay FCP RSP errors being returned to the upper SCSI layer based on the no-device-delay configuration parameter. |
| discovery-threads | Controller Specific | 1 | 1 | 32 | No | Number of ELS commands during discovery. This value specifies the number of threads permissible during device discovery. A value of 1 serializes the discovery process. |

Table 2: lpfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|------------------------|---------------------|---------|-----|-----|---------|--|
| dqfull-throttle-up-inc | Controller Specific | 1 | 0 | 128 | Yes | Amount to increment LUN queue depth each time. This parameter causes the lpfc driver to decrement a LUN's queue depth, if a queue full condition is received from the target. The queue depth will be decremented down to a minimum of 1. The variables dqfull-throttle-up-inc and dqfull-throttle-up-time are used to restore the queue depth back to the original. The dqfull-throttle-up-time parameter defines a time, in seconds, that is used to tell when to increase the current queue depth. If the current queue depth isn't equal to the lun-queue-depth, and the driver stop_send_io flag is equal to 0 for that device, increment the current queue depth by dqfull-throttle-up-inc (don't exceed the lun-queue-depth). So, if both parameters are set to 1, then driver increments the current queue depth once per second until it hits the lun-queue-depth. The only other way to restore the queue depth (besides rebooting), back to the original LUN throttle, is by running the command /usr/sbin/lpfc/resetqdepth X. This will restore the LUN throttle of all LUNs for adapter X back to the original value. |

Table 2: Ipfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|-------------------------|---------------------|---------|-----|-----|---------|--|
| dqfull-throttle-up-time | Controller Specific | 1 | 0 | 30 | Yes | Time interval (seconds) to increment LUN queue depth. Amount to increment LUN queue depth each time. This parameter causes the Ipfc driver to decrement a LUN's queue depth, if a queue full condition is received from the target. The queue depth will be decremented down to a minimum of 1. The variables dqfull-throttle-up-inc and dqfull-throttle-up-time are used to restore the queue depth back to the original. The dqfull-throttle-up-time parameter defines a time, in seconds, that is used to tell when to increase the current queue depth. If the current queue depth isn't equal to the lun-queue-depth, and the driver stop_send_io flag is equal to 0 for that device, increment the current queue depth by dqfull-throttle-up-inc (don't exceed the lun-queue-depth). So, if both parameters are set to 1, then driver increments the current queue depth once per second until it hits the lun-queue-depth. The only other way to restore the queue depth (besides rebooting), back to the original LUN throttle, is by running the command /usr/sbin/lpfc/resetqdepth X. This will restore the LUN throttle of all LUNs for adapter X back to the original value. |
| extra-io-tmo | Controller Specific | 0 | 0 | 255 | Yes | Extra timeout value, in seconds, to be applied to each FCP command sent. When connecting through a large fabric, certain devices may require a longer timeout value. |

Table 2: lpfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|---------------|---------------------|----------|-----|-----|---------|--|
| fcplib-DID | Global | Inactive | N/A | N/A | No | Setup persistent FCP bindings based on a target device's Port ID. This binding guarantees that target assignments will be preserved between reboots. The format for a bind entry is "NNNNNN:lpfcXtY" where NNNNNN is a 6 digit representation of the targets Port ID, X is the driver instance number and Y is the target assignment. Multiple entries must be separated by a comma (,) with the last entry terminated with a semi-colon (;). Target assignments, with all supported LUNs must also be configured in sd.conf, st.conf, or cmdk.conf. A sample entry follows: fcplib-DID="0000ef:lpfc0t0"; |
| fcplib-method | Controller Specific | 2 | 1 | 4 | No | Specifies the method of binding to be used. This binding method is used for persistent binding and automapped binding. A value of 1 will force WWNN binding, value of 2 will force WWPN binding and value of 3 will force DID binding. A fcplib-method value of 4 will cause target ID assignment in a private loop environment to be based on the ALPA array (hard addressed). If a binding method is not specified for a port, WWPN binding will be used. Any persistent binding whose method does not match with the bind method of the port will be ignored. A sample entry follows: lpfc0-fcplib-method=1; lpfc1-fcplib-method=2; |

Table 2: Ipfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|---------------|--------|----------|-----|-----|---------|--|
| fcp-bind-WWNN | Global | Inactive | N/A | N/A | No | <p>Setup persistent FCP bindings based on a target device's WWNN. This binding guarantees that target assignments will be preserved between reboots. The format for a bind entry is "NNNNNNNNNNNNNNNNNN:lpfcXtY" where NNNNNNNNNNNNNNNNN is a 16 digit representation of the targets WorldWide Node Name, X is the driver instance number and Y is the target assignment.</p> <p>Multiple entries must be separated by a comma (,) with the last entry terminated with a semi-colon (;). Target assignments, with all supported LUNs must also be configured in sd.conf, st.conf, or cmdk.conf. A sample entry follows:</p> <pre>fcp-bind- WWNN="20000020370c396f:l pfc1t0", "20000020370c27f7:lpfc0t2";</pre> |
| fcp-bind-WWPN | Global | Inactive | N/A | N/A | No | <p>Setup persistent FCP bindings based on a target device's WWPN. This binding guarantees that target assignments will be preserved between reboots. The format for a bind entry is "NNNNNNNNNNNNNNNNNN:lpfcXtY" where NNNNNNNNNNNNNNNNN is a 16 digit representation of the targets WorldWide Port Name, X is the driver instance number and Y is the target assignment. Multiple entries must be separated by a comma (,) with the last entry terminated with a semi-colon (;). Target assignments, with all supported LUNs must also be configured in sd.conf, st.conf, or cmdk.conf. A sample entry follows:</p> <pre>fcp-bind- WWPN="21000020370cf8263 :lpfc1t0";</pre> |

Table 2: lpfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|--------------|---------------------|---------|---------------------------------------|-----|---------|---|
| fcplib-class | Controller Specific | 3 | 2 | 3 | Yes | The lpfc driver is capable of transmitting FCP data in Class2 or Class 3. The lpfc driver defaults to using Class 3 transmission. |
| fdmi-on | Global | 0 | 0 | 2 | No | This parameter controls the fdmi capability of the lpfc driver. If set to 0 (default), fdmi is disabled. A value of 1 enables fdmi without registration of "host name" port attribute, while a value of 2 enables fdmi with registration of "host name" port attribute. |
| ip-class | Controller Specific | 3 | 2 | 2 | Yes | Fibre Channel is capable of transmitting IP data in Class2 or Class 3. The lpfc driver defaults to using Class 3 transmission. |
| link-speed | Controller Specific | 0 | 0=auto select 1=1G 2=2G 4=4G | | No | Sets link speed. |
| linkdown-tmo | Controller Specific | 30 | 0 | 255 | Yes | This variable controls how long the driver will hold I/O (0 - 255 seconds) after the link becomes inaccessible. When this timer expires, all I/O waiting to be serviced is aborted. For instance, FCP commands will be returned back to the target driver with a failure. The lower the value, the quicker the driver will fail commands back to the upper levels. There is a tradeoff here: small values risk retrying the commands when the link is bouncing; large values risk delaying the failover in a fault tolerant environment. linkdown-tmo works in conjunction with nodev-tmo. I/O will fail when either of the two timers expires. |

Table 2: lpfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|----------------------|---------------------|---------|-----|--------|---------|---|
| log-only | Controller Specific | 1 | 0 | 1 | Yes | When set to 1, log messages are only logged to syslog. When set to 0, log messages are also printed on the console. |
| log-verbose | Controller Specific | 0x0 | 0x0 | 0xffff | Yes | (bit mask) When set to non-zero this variable causes lpfc to generate additional messages concerning the state of the driver and the I/O operations it carries out. These messages may go to the system log file, /var/adm/messages and/or the system console. See Error Messages for detailed information on the bit mask. |
| lpfcNtM-lun-throttle | Controller Specific | none | 1 | 128 | No | The maximum number of outstanding commands to permit for any logical unit on a specific target. This value overrides lun-queue-depth. |
| lpfcNtM-tgt-throttle | Controller Specific | none | 1 | 10240 | No | The maximum number of outstanding commands to permit for any target, including all LUNs on that target. This value overrides tgt-queue-depth. |
| lun-queue-depth | Global | 30 | 1 | 128 | No | The driver uses this value as the default limit for the number of simultaneous commands to issue to a single logical unit on a single target on the loop. A single logical unit will never be sent more commands than allowed by lun-queue-depth; however, less may be sent when sd-max-throttle or tgt-queue-depth is reached for the entire target. |
| network-on | Controller Specific | 0 | 0 | 1 | No | This variable controls whether lpfc provides IP networking functionality over Fibre Channel. This variable is a Boolean: when zero, IP networking is disabled; when non-zero, IP networking is enabled. |

Table 2: Ipfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|-----------------|---------------------|---------|-------|------|---------|---|
| no-device-delay | Global | 1 | 0 | 30 | Yes | This variable (0 to 30 seconds) determines the length of the interval between deciding to fail an I/O because there is no way to communicate with its particular device (e.g., due to device failure or device removal) and actually failing the command. A value of zero implies no delay whatsoever. This delay is specified in seconds. A minimum value of 1 (1 second) is recommended when communicating with any Tachyon based device. |
| nodev-holdio | Controller Specific | 0 | 0=Off | 1=On | Yes | This variable controls if I/O errors are held by the driver if a FCP device on the SAN disappears. If set, I/O errors will be held until the device returns back to the SAN (potentially indefinitely). This parameter is ignored, if SCSI commands are issued in polled mode. The upper layer may retry the command once the error is returned. |
| nodev-tmo | Controller Specific | 30 | 0 | 255 | Yes | This variable (0 to 255 seconds) controls how long I/O will be held by the driver if a device on the SAN disappears. If set, I/O will be held for the specified number of seconds. If the device does not appear on the SAN before nodev-tmo seconds, then the driver will fail all held I/O and mark the device as unavailable. The upper layer may retry the command once the error is returned. |

Table 2: Ipfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|-------------|---------------------|---------|-----|-------|---------|--|
| num-bufs | Controller Specific | 128 | 64 | 4096 | No | <p>This variable specifies the number of command buffers to allocate. These buffers are used for Fibre Channel Extended Link Services (ELS), and one for each FCP command issued in SLI-2 mode. If you want to queue lots of FCP commands to the adapter, then you should increase num-bufs for better performance. These buffers consume physical memory and are also used by the device driver to process loop initialization and re-discovery activities.</p> <p>Important: The driver must always be configured with at least several dozen ELS command buffers; we recommend at least 128.</p> |
| num-iocbs | Controller Specific | 256 | 128 | 10240 | No | <p>This variable indicates the number of Input/Output control block (IOCB) buffers to allocate. IOCBs are internal data structures used to send and receive I/O requests to and from the LightPulse hardware. Too few IOCBs can temporarily prevent the driver from communicating with the adapter, thus lowering performance. (This condition is not fatal.) If you run heavy IP traffic, you should increase num-iocbs for better performance.</p> |
| post-ip-buf | Controller Specific | 128 | 64 | 1024 | No | <p>This variable specifies the number of 4K STREAMS buffers to allocate and post to the fibre channel IP ring. Increase this setting for better IP performance under heavy loading.</p> |

Table 2: lpfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|-----------------|---------------------|---------|--|-------|---------|---|
| scan-down | Controller Specific | 1 | 0=Off | 1=On | Yes | There are two scanning algorithms used to discover a node in a private loop. If scan-down is 1, devices on the private loop are scanned starting from ALPA 0x01 through ALPA 0xEF. If scan-down is 0, devices on the private loop are scanned starting from ALPA 0xEF through ALPA 0x01. Scan-down values 0 and 1 do not apply if a loop map is obtained. See the FC-AL profile for the definition of a loop map. |
| tgt-queue-depth | Global | 0 | 0 | 10240 | No | The driver uses this value as the default limit for the number of simultaneous commands to issue to a singletarget on the loop. A value of 0 causes no target throttling to occur. A single target will never be sent more commands than allowed by tgt-queue-depth; however, less may be sent when sd-max-throttle is reached for the entire target. |
| topology | Controller Specific | 0x0 | 0x0=loop , then P2P 0x2=P2P only 0x4=loop only 0x6=P2P, then loop | | No | This variable controls the Fibre Channel topology expected by lpfc at boot time. Fibre Channel offers point-to-point, fabric, and arbitrated loop topologies. To make the adapter operate as an N_Port, select point-to-point mode (used for N_Port to F_Port, and N_Port to N_Port connections). To make the adapter operate in a Fibre Channel loop as an NL_Port, select loop mode (used for private loop and public loop topologies). The driver will reject an attempt to set the topology to a value not in the above list. The auto-topology settings 0 and 6 will not work unless the adapter is using firmware version 3.20 or higher. |

Table 2: lpfc.conf Parameters (Continued)

| Parameter | Scope | Default | Min | Max | Dynamic | Comments |
|--------------|---------------------|---------|-------|--------|---------|--|
| use-adisc | Controller Specific | 0 | 0=Off | 1=On | Yes | This variable controls the ELS command used for address authentication during re-discovery upon link-up. If set, ADISC is used, otherwise, PLOGI is used. For FCP-2 devices, the driver will always use ADISC. For re-discovery due to a RSCN, the driver will always use ADISC. |
| xmt-que-size | Controller Specific | 256 | 128 | 10,240 | No | This variable specifies the number of network packets that can be queued or outstanding at any time in the driver. Increase this setting for better IP performance under heavy loading. |

Probing for FCP Disk Targets -- Configuring sd.conf

You must tell the Solaris target drivers which disks to probe for at boot time. These instructions are contained in a configuration file specific to the target driver. The disk target driver is called `sd` for Solaris 7, Solaris 8, Solaris 9 and Solaris 10. The configuration file, `sd.conf`, is usually found in `/kernel/drv`. See the `driver.conf(4)` man page for additional details.

Each disk drive on the Fibre Channel must have an entry in `sd.conf`. You may have to define additional targets and LUNs in these files. For FCP there can be up to 512 targets, with a maximum of 256 logical unit numbers (LUN) per target.

Initially, the `sd.conf` file has entries with the format:

```
name="sd" parent="lpfc" class="scsi" target=0 lun=0
```

The `class` keyword ensures that Solaris specifically probes all adapters controlled by all drivers that register themselves as `class="scsi"`, for the specified targets and LUNs. The `lpfc` driver registers itself as `class="scsi"`.

The installation procedure for `lpfc` that is automatically added to `sd.conf` follows this format:

```
#name="sd" parent="lpfc" target=16 lun=0;
#name="sd" parent="lpfc" target=17 lun=0;
#name="sd" parent="lpfc" target=17 lun=1;
#name="sd" parent="lpfc" target=17 lun=2;
#name="sd" parent="lpfc" target=17 lun=3;
```

If you want to use disks (for example, RAID arrays) with non-zero LUNs, or with more than 18 targets, simply add lines to the target driver configuration file specifying the desired target ID and LUN(s). For instance, you can include the following example in your target driver configuration file:

```
name="sd" parent="lpfc" target=17 lun=0;
name="sd" parent="lpfc" target=17 lun=1;
```

or

```
name="sd" class="scsi" target=17 lun=0;  
name="sd" class="scsi" target=17 lun=1;
```

The parent keyword ensures that Solaris specifically probes all adapters controlled by the lpfc driver for the specified targets and LUNs.

Note: Target probing changes in sd.conf will take effect only after the system has been rebooted.

Probing for FCP Tape Targets -- Configuring st.conf

You must tell the Solaris target drivers which tapes to probe for at boot time. These instructions are contained in a configuration file specific to the target driver. The configuration file, st.conf, is usually found in /kernel/drv. See the driver.conf(4) man page for additional details. The disk target driver is called st, and entries must be included in this file in order to run tape devices.

No entries are automatically added to the st.conf file. You can include the following example in your target driver configuration file:

```
name="st" parent="lpfc" target=17 lun=0;  
name="st" parent="lpfc" target=17 lun=1;
```

or

```
name="st" class="scsi" target=17 lun=0;  
name="st" class="scsi" target=17 lun=1;
```

The parent keyword ensures that Solaris specifically probes all adapters controlled by the lpfc driver for the specified targets and LUNs.

Note: Target probing changes in st.conf will take effect only after the system has been rebooted.


Along with configuring the target and LUNs, the tape-config-list section of the st.conf file must be configured for Solaris to correctly communicate with the tape drives. Contact your tape vendor for assistance.

Viewing Target Mapping and Set Up Persistent Binding

The **Target Mapping** tab in HBAnyware enables you to view current target mapping and to set up persistent binding. You can also set up persistent binding using lputil.

Viewing Target Mapping Using HBAnyware

To view the **Target Mapping** tab:

1. Start HBAnyware.
2. Do one of the following:
 - From the menu bar, click **View**, then click **Group HBAs by Host Name**.
 - From the toolbar, click the **Group HBAs by Host Name** button. 
3. Click a target in the discovery tree.
4. Click on the **Target Mapping** tab.

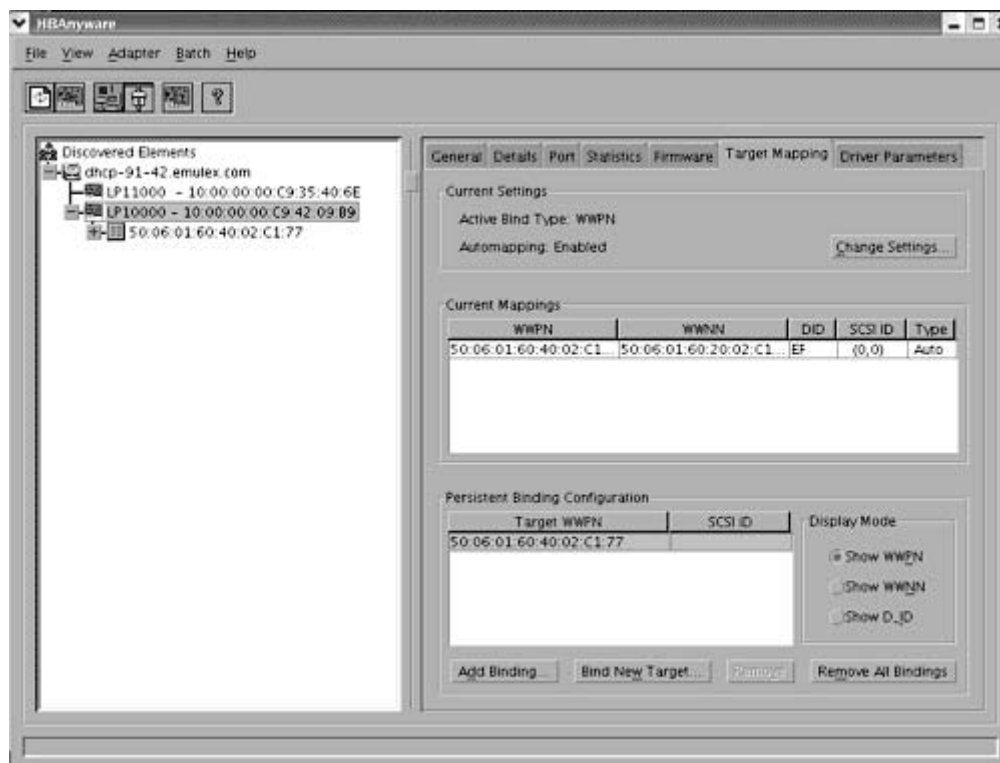


Figure 26: The HBAnyware Target Mapping Tab

Setting Up Persistent Binding Using HBAnyware

When you create a persistent binding, HBAnyware tries to make that binding dynamic. However, the binding must meet all of the following criteria to be dynamic:

- The SCSI ID (target/bus combination) specified in the binding request must not be mapped to another target. For example, the SCSI ID must not already appear in the 'Current Mappings' table under 'SCSI ID'. If the SCSI ID is already in use, then the binding cannot be made dynamic, and a reboot is required.

- The target (WWPN, WWNN or DID) specified in the binding request must not be mapped to a SCSI ID. If the desired target is already mapped, then a reboot is required.
- The 'Bind Type Selection' (WWPN, WWNN or DID) specified in the binding request must match the currently active bind type shown under “Current Settings” section of the **Target Mapping** tab. If they do not match, then the binding cannot be made active.

To set up persistent binding:

1. Start HBAnyware.
2. In the directory tree, click the host bus adapter (HBA) for which you want to set up persistent binding.
3. Click the **Target Mapping** tab. All targets are displayed.
4. The information for each currently defined mapping includes the world wide port name (WWPN), world wide node name (WWNN), device ID (D_ID), SCSI ID, or Bind Type. The type can be either 'PB', indicating that the mapping was the result of a persistent binding, or 'Auto', indicating that the target was automapped. In the Display Mode section, choose the display mode you want to use.
5. If you want to change the Active Bind Type (the mode used to persistently bind target mappings) or Automapping setting, click **Change Settings**. Select the Bind Type (WWPN, WWNN or D_ID), and set Automapping to Enabled or Disabled.

Note: All mapped targets, whether automapped or resulting from a persistent binding configuration, will have entries in the “Current Mappings” table on the **Target Mapping** dialog box.

If the binding that you defined has been successfully activated, you will see the following message:

“The new binding has been created and is currently active.”

If, however, the binding was successfully created, but could not be made active, you will see the following message:

“The new binding has been created. Note that this binding will not become active until after you have rebooted the system.” Generally, you should ensure that the bind type in the Current Settings section of the **Target Mapping** dialog box is the same as the type of binding selected in the Persistent Binding Configuration section of the dialog box.

To add a persistent binding:

1. In the Targets Table, click the target that you want to bind.

2. Click **Add Binding**. The **Add Persistent Binding** dialog box is displayed.

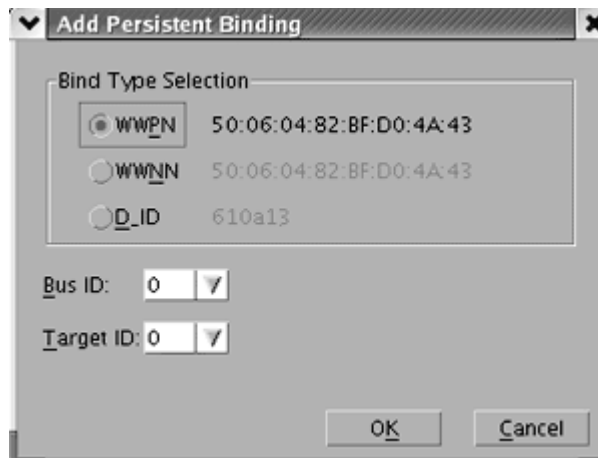


Figure 27: The HBAnyware Add Persistent Binding Dialog Box

3. Select the Bind Type that you want to use (WWPN, WWNN or D_ID).
4. Select the Bus ID and Target ID that you want to bind, and click **OK**.

Note: All mapped targets, whether automapped or resulting from a persistent binding configuration, will have entries in the “Current Mappings” table on the **Target Mapping** dialog box.

If the binding that you defined has been successfully activated, you will see the following message:

“The new binding has been created and is currently active.”

If, however, the binding was successfully created, but could not be made active, you will see the following message:

“The new binding has been created. Note that this binding will not become active until after you have rebooted the system.” Generally, you should ensure that the bind type in the Current Settings section of the **Target Mapping** dialog box is the same as the type of binding selected in the Persistent Binding Configuration section of the dialog box.

To bind a target that does not appear in the Persistent Binding Table:

1. Click **Bind New Target**. The **Bind New Target** dialog box is displayed.




Figure 28: The Bind New Target Dialog Box

2. Click the type of binding you want to use, and type the WWPN, WWNN or D_ID you want to bind to the target.
3. Select the Bus ID and Target ID that you want to bind, and click **OK**.

Note: A target will not appear on the target list if automapping has been disabled and the target is not already persistently bound.

Adding New Targets Using sd.conf (Solaris 8)

You can perform on-the-fly configuration changes, without rebooting, using HBAnyware. For Solaris 8, you must first add the new targets to the sd.conf file.

To add new targets using sd.conf (Solaris 8):

1. Edit the Solaris SCSI configuration file (sd.conf):

```
#vi /kernel/drv/sd.conf
.
.
.
name="sd" parent="lpfc" target=17 lun=1;
name="sd" parent="lpfc" target=18 lun=10;
name="sd" parent="lpfc" target=19 lun=15;
.
.
.
```

2. Save the file and exit vi.

Changing Parameters or Bindings

To change parameters or bindings in Solaris 9, edit the sd.conf file as shown above and force a reread of the file with the `update_drv -f sd` command.

To change parameters or bindings in Solaris 7 or 8:

1. Quiesce all I/O on the device.
2. Unconfigure all ports with open instances to the driver.
3. Unload the driver using the `modunload` command. (See "Loading or Unloading the Driver without Rebooting" on page 31 for more information.)
4. Reload the driver using the `modload` command. (See "Loading or Unloading the Driver without Rebooting" on page 31 for more information.)

Setting Up Persistent Binding using lputil

Persistent binding allows you to permanently assign a system SCSI target ID to a specific Fibre Channel (FC) device even though the device's ID on the FC loop (D_ID) may be different each time the FC loop initializes. This capability is useful in multi-server environments that share a device. You can simplify system management by having multiple servers use the same SCSI target ID when referring to the shared FC device.

To set up persistent binding using lputil:

1. Start the lputil utility, type:

```
/usr/sbin/lpfc/lputil
```

The Main menu opens:

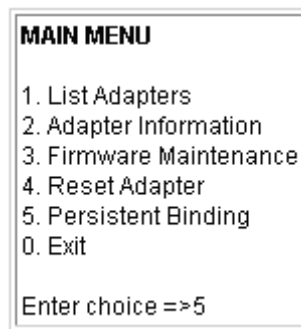


Figure 29: The *lputil*, Main menu

2. Select choice #5. The Persistent Binding menu opens:

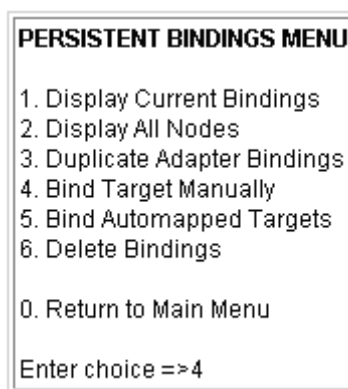


Figure 30: The *lputil*, Persistent Bindings menu

3. To establish new bindings, select option => 4.
4. Select an HBA.
5. Selecting a binding method.

Note: In a fabric environment, the D-ID may change when the system is rebooted. We suggest binding to the Node Name or Port Name in a fabric environment.

1. By Node Name
2. By Port Name
3. By D_ID
0. Cancel
6. Enter the Node Name, Port Name, or D_ID, depending on binding method.
7. Enter the target number => 0.

Setting Up Target/LUN Blocking Using *sd.conf*

The class keyword ("scsi") ensures that Solaris specifically probes all adapters controlled by all driver that register themselves as class="scsi". The parent keyword ("lpfc") ensures that Solaris specifically probes all adapters controlled by the lpfc driver for the specified targets and LUNS. The class and parent keywords cause the SCSI layer to probe multiple adapters, even multiple adapters across multiple drivers. This method limits the SCSI layer probing of targets and LUNs on an adapter-by-adapter basis. This gives you control over which targets and LUNs are seen by each initiator (target/LUN blocking).

To set up target/LUN blocking using sd.conf:

1. Reboot the system with the adapter installed.
2. Check the output of `dmesg(1M)`. This message displays in the following format:

```
NOTICE: Device Path for interface lpfcX:/pci@1f.0/pci@1/fibre-  
channel@3
```

Note: Where `lpfcX` is the interface for a specific adapter and `/pci@1f.0/pci@1/fibre-channel@3` is the device path for the specific adapter.

3. Add entries to the `sd.conf` file in the following format:

```
name="sd" parent="lpfc" target=16 lun=0 hba="lpfcX";
```

Note: This entry does not cancel the effect of any other `parent="lpfc"` or `class="scsi"` entries for `target=16 lun=0`. If you want the SCSI layer to probe only for `target=16 lun=0` on device `lpfcX`, the `parent="lpfc"` or `class="scsi"` entries for `target=16 lun=0` need to be deleted. You can cause system problems if certain `class="scsi"` entries are deleted. These entries are used by the SCSI adapter, so if there is a SCSI boot disk at `target=0 lun=0` whose probe entry has been deleted, the system won't boot. Similarly, if any SCSI target's probe entry is deleted, that device won't operate. To guarantee that the Fibre Channel and SCSI probing won't conflict, use persistent binding to assign FC devices target numbers greater than 15. Persistent binding can be used to perform target blocking but not LUN blocking.

Enabling IP Networking

Overview

Usually, IP networking is enabled during driver installation. Use this procedure to enable IP networking if it was not enabled during driver installation. In addition, if there is more than one adapter in the host, and both storage and IP networking functions are desired, this procedure will tell the system which adapters will have IP functionality. All others will have storage functionality.

To enable IP networking:

1. If IP networking support was not enabled during initial driver installation, update the `lpfc.conf` file to enable IP networking.
2. Modify the `/etc/hosts` file to add the hostname and IP address of the Solaris host.
3. Create a `hostname.lpfcX` file and assign the network name for additional host adapter X.
4. Reboot the system. Upon rebooting, Solaris will automatically configure and set up the additional adapters for networking.
5. Verify the network connection by pinging a known IP address.

Enabling the Networking Driver Parameter

IP networking defaults to disabled and is usually enabled during driver installation. If you are enabling IP networking after the driver is installed, you need to set the network-on driver parameter to 1.

To enable the networking driver parameter:

1. Use any editor (for example, `vi`) to update `lpfc.conf` and enable networking:

```
vi /kernel/drv/lpfc.conf
```
2. To enable IP networking, enter a binary 1 for the network-on driver parameter (1 = on, 0 = off).

Creating Hostname.lpfc# Files

Required Hostname File

When you install the driver for the first time, you are required to create a file named `hostname.lpfc0` for the first host bus adapter (HBA) found. This file corresponds to the device instance number. If for any reason the device instance number changes, the file name must change accordingly. Use any text editor to edit the `etc/hosts` file, to associate an IP address with the HBA's network host name.

Additional Hostname Files

You must also manually create a `hostname.lpfc#` file for each additional HBA in the system. Each `hostname.lpfc#` file contains the network host name of the HBA. Usually you create these additional hostnames during the installation process, however you can also add HBAs to your system after you install the driver.

Use any text editor to edit the `/etc/hostname.lpfcx` (where x is the HBA number in decimal format).

Example

```
123.456.7.8 fibre1
123.456.8.9 fibre2
123.456.8.1 myhost
```

To create a `hostname.lpfc#` file:

1. Using any text editor, create a new file with the name 'fibre1' in it. Save this file as `hostname.lpfc0`. This is the required hostname file.
2. Create another new file with the name 'fibre2' in it. Save this file as `hostname.lpfc1`.
3. Create another new file with the name 'myhost' in it. Save this file as `hostname.lpfc2`.
4. If you are creating additional hostname files during an installation, continue with step 13 on the Install a Driver for the First Time page in the Installation manual.

Note: If you are adding an HBA to your network after the driver is installed, a reboot is necessary to configure and set up the additional HBA for networking.

Verifying the Network Connection

Verify the connection to the network by pinging a known IP address.

```
ping ip_address
```

At this point, you may use the host adapter or configure additional Solaris parameters to increase performance or optimize driver operation in a unique environment. To perform additional configuration changes, consult your Solaris documentation.

Enabling IP Networking for Multiple Adapters

To enable networking on two or more installed host adapters in the same system, you must create additional `hostname.lpfcX` files manually (where X is the instance number of the host adapter being configured). Upon rebooting, Solaris will automatically configure and set up the additional adapters for networking.

If IP networking support is enabled after the initial driver installation, the network-on setting must be updated in `lpfc.conf`.

Any name given in the `hostname.lpfcX` files must have a valid IP address associated with it. The IP address is established through NIS, DNS or another addressing scheme.

Alternately, rerun the installation procedure (doing so is called an update). Before you update the driver, be sure to save the hostname file and IP address information in a temp file. Running an update install provides the opportunity to enable IP networking support, then restore all hostname file and IP address information, as well as all saved configuration file settings. Reboot the system for these changes to take effect. See the online Installation manual for more information about updating the driver.

Creating Additional Hostname.lpfc# Files

When the driver is installed, a file named `hostname.lpfc0` is created for the first adapter found. (This file corresponds to the device instance number. If for any reason the device instance number changes, the file name must change accordingly.) You must manually create a `hostname.lpfc#` file for each additional adapter in the system. Each `hostname.lpfc#` file contains the network host name of the adapter.

For example, if we have three adapters:

```
123.456.7.8 fibre1
123.456.8.9 fibre2
123.456.8.1 myhost
```

the file `hostname.lpfc0` is created and contains the name 'fibre1'.

Using any text editor, create a new file with the name 'fibre2' in it. Save this file as hostname.lpfc1.

Again, create another file with the name 'myhost' in it. Save this file as hostname.lpfc2.

Configuring the System-Wide File (/etc/system)

Normally, you do not need to modify the contents of /etc/system. Most required changes for /etc/system are made automatically during the installation. All of the commonly used parameters are in lpfc.conf and the target driver configuration file. Listed below are a few cases that may require modification of the /etc/system.

Driver Installed in a Non-Standard Directory (moddir)

The moddir parameter specifies a series of directories that the kernel will search for loadable device drivers at boot time. If you specify a non-standard installation directory, the installation process will create or update the moddir parameter to include the new directory. If this fails, or if you later decide to move the lpfc.conf file, you must create or update moddir.

IP Networking Enabled (forceload)

This parameter loads the specified modules at boot time, just before mounting the root filesystem, rather than at first reference. You will find this example in /etc/system after installing the driver (assuming you enabled IP networking):

```
#forceload: drv/clone
```

The kernel locates the modules specified in forceload by consulting moddir. If you enable IP networking for the host adapter driver, the parameter forceload must be set to drv/clone for correct operation. The lpfn driver is dependent on the clone driver. Normally, the clone driver would load itself automatically when the driver calls it the first time. However, the lpfn driver is loaded so early in the system boot sequence that a forced loading of the clone driver is required.

HBAnyware Security

Introduction

After HBAnyware, which includes the HBAnyware utility and remote server, is installed on a group of systems, HBAnyware can remotely access and manage the HBAs on any systems in the group. This may not be a desirable situation, because any system can perform actions such as resetting boards or downloading firmware.

The HBAnyware security package can be used to control which HBAnyware systems can remotely access and manage HBAs on other systems in a Fibre Channel network. HBAnyware security is systems-based, not user-based. Anyone with access to a system that has been granted HBAnyware client access to remote HBAs can manage those HBAs. Any unsecured system is still remotely accessible by the HBAnyware client software (HBAnyware utility).

The HBAnyware security software is designed to provide two main security features:

1. Prevent remote HBA management from systems that the administrator does not want to have this capability.
2. Prevent an accidental operation (such as firmware download) on a remote HBA. In this case, the administrator does not want to have access to HBAs in systems he or she is not responsible for maintaining.

The first time the HBAnyware Security Configurator is run on a system in an environment where no security has been configured, the initial Access Control Group (ACG) is created. At this point, only this system has remote access to the HBAs in the systems in the ACG. They are no longer remotely accessible from any other system.

Subsequently, additional Access Sub-Groups (ASGs) can be created. This grants systems in the ACG the ability to remotely access the HBAs of other selected systems in the ACG.

Starting the HBAnyware Security Configurator

Prerequisites

Before you can start the HBAnyware Security Configurator, you must have the following items installed on your system. See the online Installation manual for more information.

- The Emulex driver for Solaris
- The HBAnyware and lputil Utilities
- The HBAnyware Security Configurator

Note: Before you start the Configurator, you must make sure that all of the systems that are part of, or will be part of, the security configuration are online on the Fibre Channel network so that they receive updates or changes made to the security configuration.

Any system that is already part of the security installation might not run with the proper security attributes, if updates to the security configuration are made while it is offline.

Any system that is part of the security installation and that is offline when the HBAnyware Security Configurator starts will not be available for security configuration changes even if it is brought online while the Configurator is running.

Procedure

To start the HBAware Security Configurator:

1. Run the `/usr/sbin/hbanyware/ssc` script. Type:

```
/usr/sbin/hbanyware/ssc
```

Running the Configurator for the First Time/Creating the ACG

When the HBAware Security software is installed on a system and the HBAware Security Configurator is run for the first time, that system becomes the Master Security Client (MSC). All of the available servers are discovered and available to become part of the system Access Control Group (ACG). You select the systems to be added to the ACG, and the security configuration is updated on all of the selected servers as well as on the initial system. This selection constitutes the participating platforms in this security installation.

To create the ACG:

1. Start the HBAware Security Configurator for the first time in an unsecure environment. The computer from which you run the Configurator will become the MSC. The following message is displayed:




Figure 31: The HBAware Security Configurator "Unsecure System" message

2. Click **OK**. The **Access Control Group** tab is displayed.

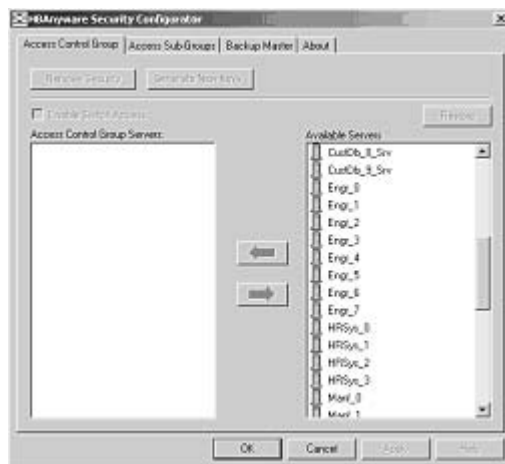


Figure 32: The Access Control Group Tab

3. Select the unsecured servers that you want to add to the ACG from the Available Servers list.




Figure 33: The Access Group Control Tab

4. Click the **left arrow** to add the servers to the Access Control Group Servers list.
5. Click **OK** or **Apply**.

Designating a Master Security Client

The first time you run the HBAAnyware Security Configurator on any system in a Fibre Channel network, that system becomes the MSC (Master Security Client). See “Running the Configurator for the First Time” on page 85 for more information.

Access Control Groups

Introduction

The **Access Control Group** tab shows the systems that are part of a client's Access Control Group (ACG) and, from the Master Security Client (MSC), allows you to select the systems that belong to the ACG.

Access Control Group Tab on the MSC

On the MSC, you select or deselect the systems that are to be part of the security installation in the **Access Control Group** tab. When you select unsecure systems and move them to the Access Control Group Servers list, these systems are updated to secure them and bring them into the MSC's ACG.

When you select systems in the ACG and move them to the Available Servers list, the security configuration for those systems is updated to make them unsecure. After you have configured security from the MSC for the first time, the **Access Control Group** tab looks similar to the following:

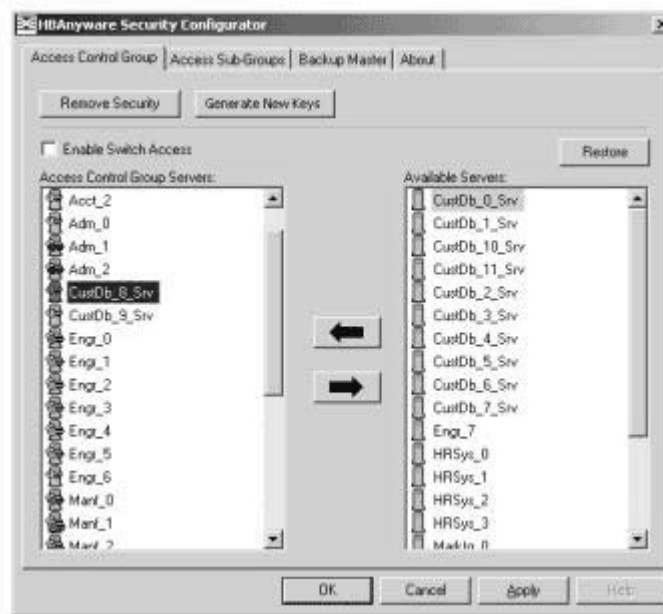


Figure 34: The Access Control Group Tab on a non-MSC system

Access Control Group Tab on a Non-MSC

On a non-MSC system, the **Access Control Group** tab shows the systems that are part of the client's ACG. You cannot modify the ACG on a non-MSC. (You can modify the ACG only on the MSC or a client higher in the security topology's hierarchy.) The **ACG** tab on a non-MSC system looks similar to the following:







Figure 35: The Access Control Group Tab after MSC security is configured


ACG Icons


Depending on the configured security topology, a system can be a server in one or more ACGs. It can also be a client to an ACG. The following icons indicate the state of each of the systems in the Access Control Group Servers list.

- 

The system is a secure server in the ACG. It does not belong to an Access Sub-Group (ASG). You can remove this system from the ACG.
- 

The system is a secure server in the ACG and belongs to one or more ASGs. You can remove this system from the ACG.
- 

The system is a secure server in the ACG and a client to an ASG. You cannot remove this system from the ACG until you remove it as a client from the ASG.
- 

The system is a secure server in the ACG, a secure server in one or more ASGs and a client to an ASG. You cannot remove this system from the ACG until you remove it as a client from the ASGs.
- 

The system is a Backup Master. You cannot remove this system from the ACG until you remove it as a Backup Master.

Creating the ACG

When the HBAnyware Security software is installed on a system and the HBAnyware Security Configurator is run for the first time, that system becomes the Master Security Client (MSC). All of the available servers are discovered and available to become part of the system Access Control Group (ACG). You select the systems to be added to the ACG, and the security configuration is updated on all of the selected servers as well as on the initial system. This selection constitutes the participating platforms in this security installation.

To create the ACG:

1. Start the HBAnyware Security Configurator for the first time in an unsecure environment. The computer from which you run the Configurator will become the MSC. The following message is displayed:




Figure 36: HBAnyware Security Configurator message

2. Select the unsecured servers that you want to add to the ACG from the Available Servers list.

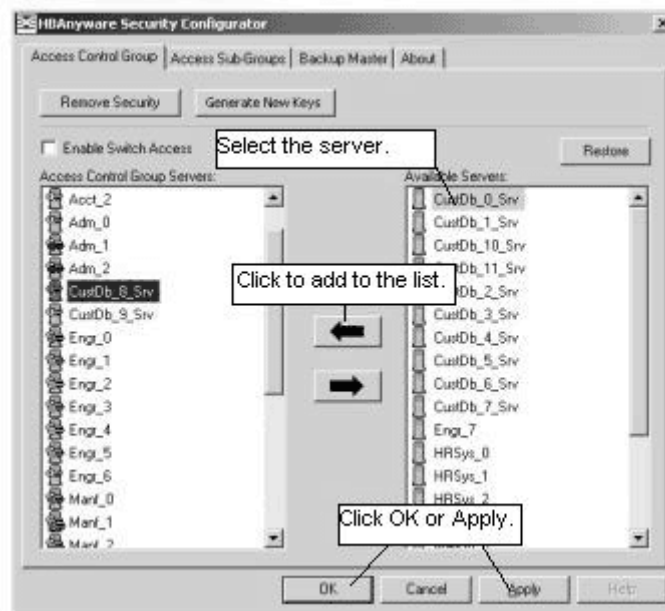


Figure 37: The Access Group Control Tab

3. Click the **left arrow** to add the servers to the Access Control Group Servers list.
4. Click **OK** or **Apply**.

Adding a Server to the ACG

After you create the initial Access Control Group (ACG) on the Master Security Client (MSC), you may want to add unsecured servers to the ACG.

To add servers to the ACG:

1. Start the HBAAnyware Security Configurator.

2. On the **Access Control Group** tab, from the Available Servers list, select the unsecured servers that you want to add to the ACG.

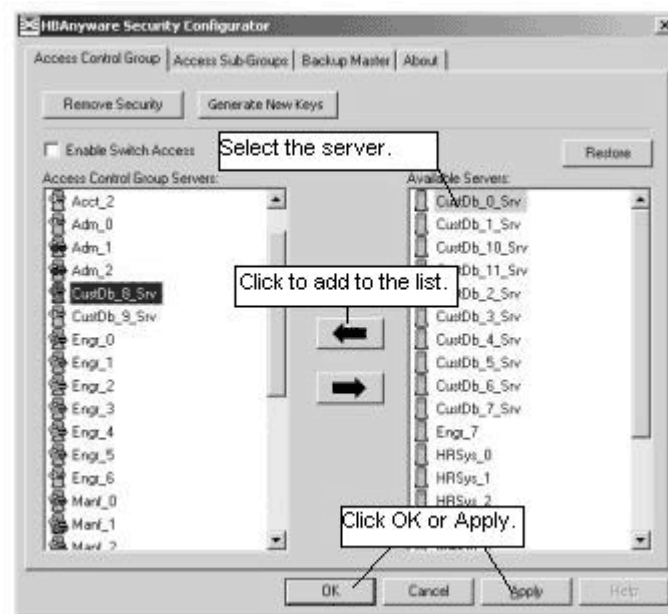


Figure 38: The Access Control Group Tab

3. Click the **left arrow** to add the server to the Access Control Group Servers list.
4. Click **OK** or **Apply**.

Deleting a Server from the ACG

To delete a server from the Access Control Group (ACG):

1. Start the HBAware Security Configurator.

- On the **Access Control Group** tab, from the Access Control Group Servers list, select the secured systems that you want to delete from the ACG.

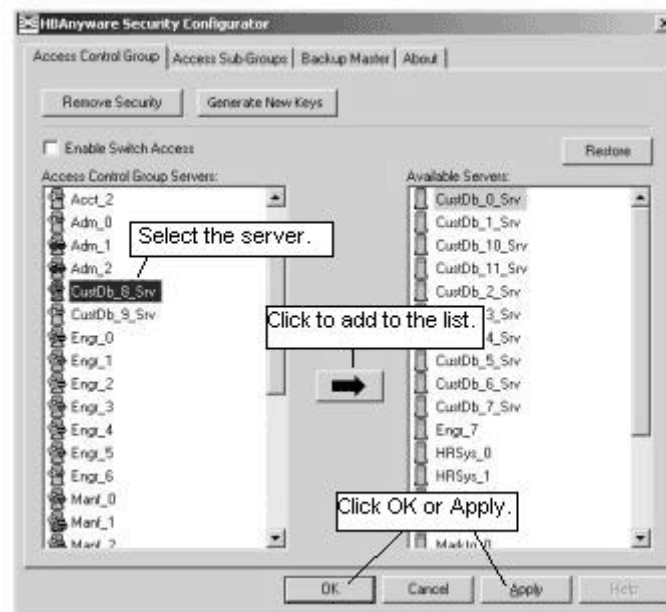


Figure 39: The Access Control Group Tab

- Click the **right arrow** to remove the servers from the Access Control Group Servers list.
- Click **OK** or **Apply**.

Removing Security from all Servers in the ACG

You can remove security from all systems only from the Master Security Client (MSC). Removing the entire security topology on all of the servers in the MSC's ACG puts the servers in an unsecured state. The MSC is also put in an unsecured state; consequently, it is no longer the MSC. Any participating systems that are not online will not receive the 'remove security' configuration update, and as a result will no longer be accessible remotely.

To remove security from all servers in the ACG:

1. Start the HBAnyware Security Configurator. The **Access Control Group** tab is displayed.

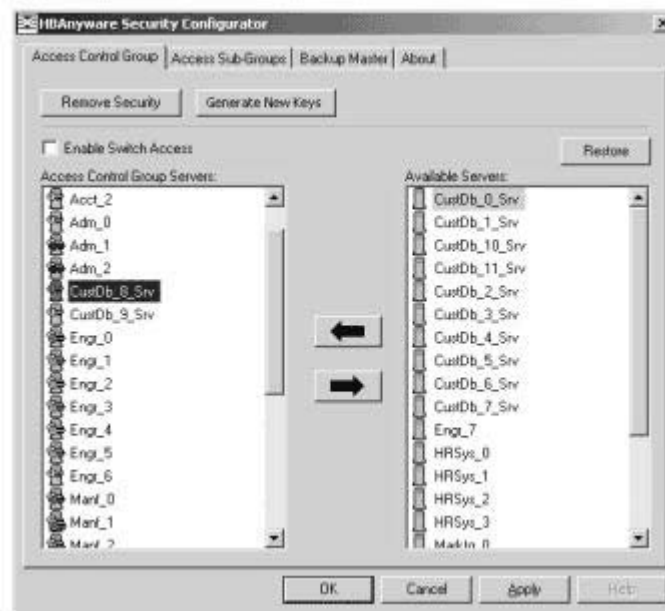


Figure 40: The Access Control Group Tab

2. On the **Access Control Group** tab, click the **Remove Security** button. The following message is displayed:



Figure 41: The HBAnyware Security Configurator "Warning" Dialog Box

3. Click **Yes**. Security is removed from all servers in the ACG.

Generating New Security Keys

You can generate new security keys only from a Master Security Client (MSC). After the new security keys are generated, they are automatically sent to all of the remote servers in the Access Control Group (ACG).

Note: All the servers that are part of the ACG must be online when this procedure is performed so that they may receive the new keys. Any servers that do not receive the new keys will no longer be accessible remotely.

To generate new security keys for all servers in the ACG:

1. From the MSC, start the HBAAnyware Security Configurator. The **Access Control Group** tab is displayed.

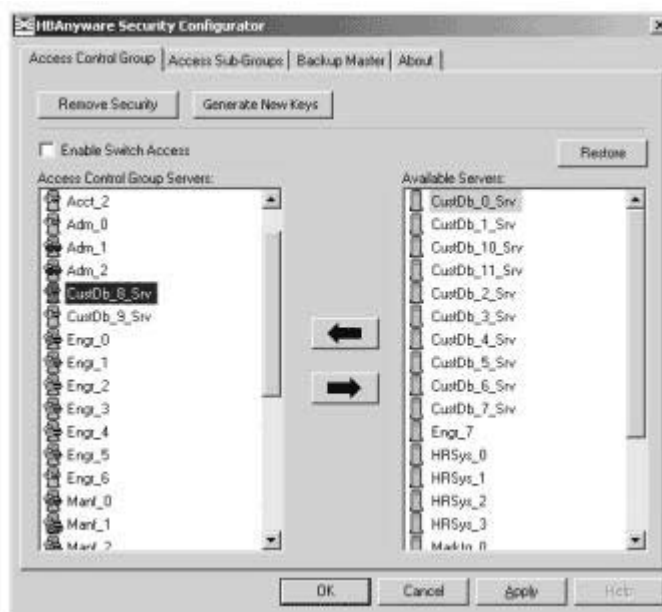


Figure 42: The Access Control Group Tab

2. On the **Access Control Group** tab, click the **Generate New Keys** button. A dialog box warns you that you are about to generate new security keys for all systems.
3. Click **Yes**. The new keys are generated and sent to all of the remote servers in the ACG.

Restoring the ACG to Its Last Saved Configuration

You can restore the ACG to its last saved configuration, if there are unsaved changes to the ACG, only from the Master Security Client (MSC).

To restore the ACG to its last saved configuration:

1. From the **Access Control Group** tab on the MSC, click the **Restore** button.

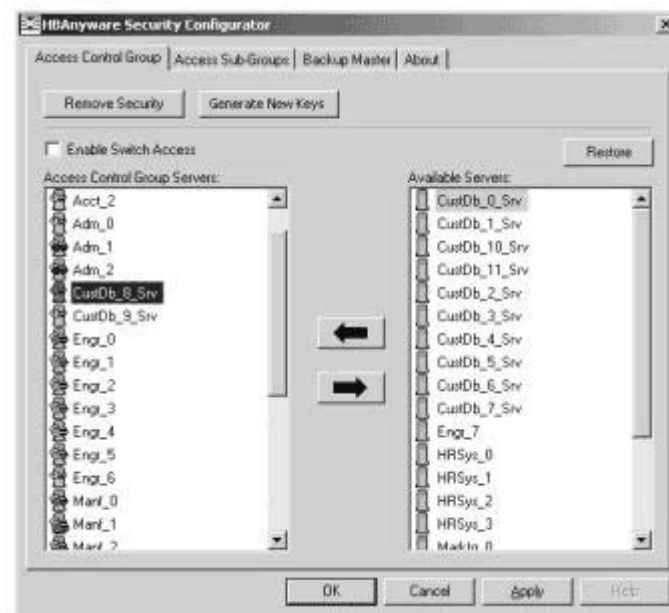


Figure 43: The Access Control Group Tab

Accessing a Switch

You can enable switch access only on a Master Security Client (MSC). Switch access grants the client access rights to a switch to remotely access HBAs on servers in the Access Control Group (ACG).

To enable switch access:

1. Start the HBAAnyware Security Configurator.

- From the **Access Control Group** tab, check **Enable Switch Access**.

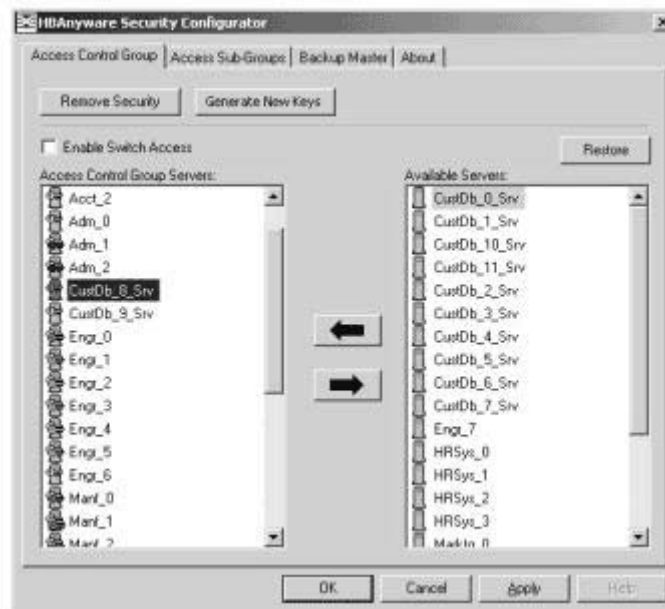


Figure 44: The Access Control Group Tab

Accessing Sub-Groups

Introduction

The **Access Sub-Group** tab allows you to create multiple Access Sub-Groups (ASGs) and multiple levels (tiers) in the security topology hierarchy. The hierarchy can be as many levels deep as desired. However, it is recommended the hierarchy extend no more than three levels deep, as it becomes increasingly difficult to keep track of the topology the deeper it goes. The hierarchy of ASGs is displayed in the **Access Sub-Groups** tab as a tree. You can create, modify and delete ASGs at each level in this tree.

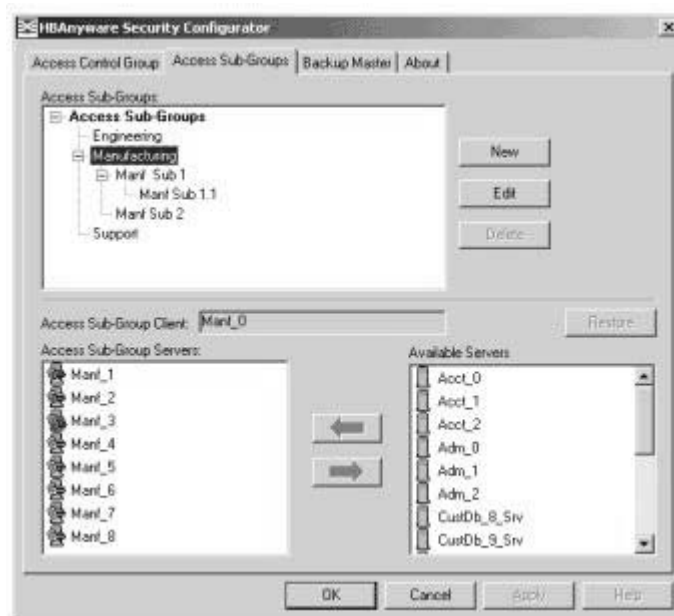


Figure 45: The Access Sub-Groups Tab

ASG Icons

The following icons indicate the state of each of the servers in the Access Sub-Group Servers list.



The system is a server in the ASG but not in any child ASGs. You can remove it from the ASG.



The system is a server in the ASG and at least one child ASG. You cannot remove it from the ASG until you remove it from the child ASGs.



The system is a server in the ASG and a client to a child ASG. You cannot remove it from the ASG until you remove it as a client from the child ASG (by either deleting or editing the child ASG).



The system is a server in the ASG, a server in at least one other child ASG and a client to a child ASG. You cannot remove it from the ASG until you remove it from the child ASGs and as a client from the child ASG (by either deleting or editing the child ASG).



The system is a server in the ASG and a client to a non-child ASG. You can remove it from the ASG.



The system is a server in the ASG, a server in at least one child ASG, and a client to a non-child ASG. You cannot remove it from the ASG until you remove it from the child ASGs.

Creating an ASG

You create a new Access Sub-Group (ASG) by selecting one system from the Access Control Group (ACG) to be the client, and some or all of the other systems to be servers to this client, thus defining the new client's ACG. When the HBAAnyware Security Configurator is run on the new client, the displayed ACG shows the servers that were configured in the ASG by its parent client.

To create an ASG:

1. Start the HBAnyware Security Configurator.
2. Click the **Access Sub-Groups** tab.




Figure 46: The Access Sub-Groups Tab

3. Click **New**. The **New Access Sub-Group** dialog box is displayed.

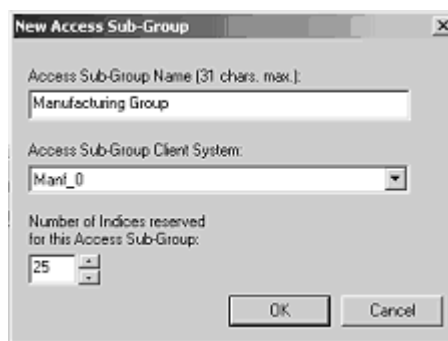


Figure 47: The New Access Sub-Group Tab

4. Enter the ASG information:
 - Access Sub-Group Name: Enter the name of the ASG. The ASG name is for identification purposes only. It does not provide any security function. Provide a name that will make it easy to remember the systems that are part of the ASG. The name can contain any alphanumeric characters, symbols or spaces (up to 31). At each level of the security topology, each ASG name must be unique. If the name is not unique at its level, an error message informs you of this when you click **OK**.
 - Access Sub-Group Client System: Select the system that is to be the client.
 - Number of indices reserved for this Access Sub-Group: Select the number of 'indices' you want to reserve for the client system of the new ASG. This number reflects the number of subsequent 'child' ASGs that can subsequently be created on the new client's system. See the Reserved Indices topic (under Access Sub-Groups in this manual) for examples.
5. Click **OK** in the **New Access Sub-Group** dialog box. The ASG is created.

Reserved Indices - Examples

A particular security installation can support the creation of several hundred access groups (ACGs and ASGs). When you create each new access group, you allocate some number of 'indices' to the client system of the new ASG. This number reflects the number of subsequent 'child' ASGs that can subsequently be created at the new client's system.

- If zero indices are reserved, you cannot create any lower-level ASG under the client of the new ASG. Thus, for example, if you want to implement a multi-tiered security architecture consisting of many ASGs, and you wanted to create them all from the Master Security Client (MSC), zero indices would be allocated to each of the new ASGs client platforms when they are created.
- If you create an ASG, and you reserve 25 indices for the new ASG client platform, a child ASG created by this platform will have a maximum of only 24 indices available to be reserved (one is taken by the creation of the child ASG itself). This continues down the ASG hierarchy as each lower level ASG is created.
- When you create an ASG from the MSC, a maximum of 50 indices (or less if fewer are available) can be reserved. For all other clients, the maximum depends on how many indices were reserved to that client when its ASG was created, and on how many it has subsequently allocated to its ASGs.

Adding a Server to an ASG

To add a server to an ASG:

1. Start the HBAnyware Security Configurator.
2. Click the **Access Sub-Group** tab.




Figure 48: The Access Sub-Group Tab

3. The name of the ASG is displayed in the Access Sub-Groups tree. From the Available Servers list, select the servers to be added to the ASG.
4. Click the **left arrow** to move the servers to the Access Sub-Group Servers list.
5. Click **OK** or **Apply** to update servers, adding them to the ASG. The new client can remotely manage the HBAs on those servers using the HBAnyware utility.

Deleting an ASG

Only a leaf node ASG may be deleted (i.e. not ASGs underneath it in the tree). If an ASG has at least one child ASG, those child ASGs must be deleted first.

To delete an ASG:

1. From the Access Sub-Group tree, select the leaf node ASG you wish to delete.
2. Click the **Delete** button. A dialog box appears warning you that if you continue the access sub-group will be deleted.
3. Click **Yes**. This operation is immediate. There is no need to click the **OK** or **Apply** button under the tab.

Restoring an ASG to Its Last Saved Configuration

You can restore an Access Sub-Group (ASG) to its last saved configuration if there are unsaved changes to it.

To restore an ASG to its last saved configuration:

1. Click the **Access Sub-Group** tab.



Figure 49: The Access Sub-Groups Tab

2. Select the ASG whose configuration you want to restore.
3. Click **Restore**.
4. Click **OK** or **Apply** to save your changes.

Editing an ASG

You can change the name, client system or reserved indices of an Access Sub-Group (ASG).

To edit an ASG:

1. Start the HBAAnyware Security Configurator.

2. Click the **Access Sub-Group** tab.

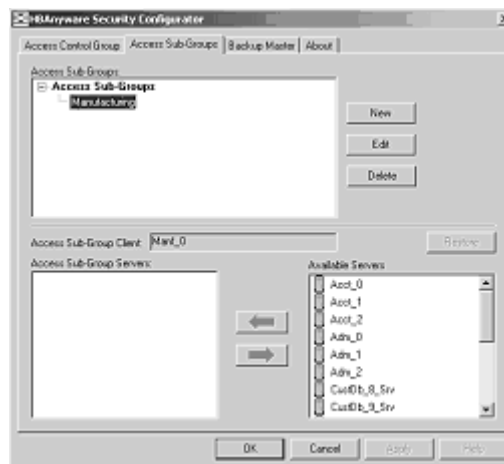


Figure 50: The Access Sub-Groups Tab

3. Select the ASG you want to edit.
4. Click **Edit**. The **Edit Access Sub-Group** dialog box is displayed.

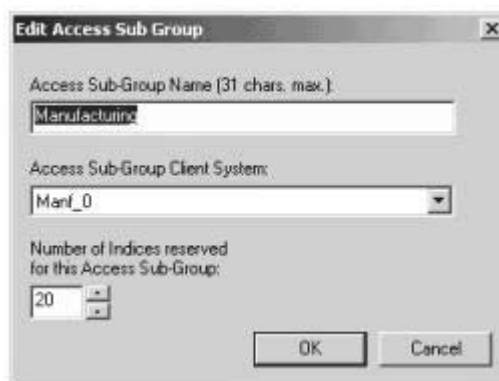


Figure 51: The Edit Access Sub Group Dialog Box

5. Change the ASG information:
 - **Access Sub-Group Name:** Change the name of the ASG. The ASG name is for identification purposes only. It does not provide any security function. Provide a name that will make it easy to remember the systems that are part of the ASG. The name can contain any alphanumeric characters, symbols or spaces (up to 31). At each level of the security topology, each ASG name must be unique. If the name is not unique at its level, an error message informs you of this when you click **OK**.
 - **Access Sub-Group Client System:** Select the new system that is to be the client. If the Configurator is running on a system connected to more than one fabric, the client list contains only those systems that can be accessed by the original client of the ASG.
 - **Number of indices reserved for this Access Sub-Group:** Select the new number of 'indices' you want to reserve for the client system of the new ASG. This number reflects the number of subsequent 'child' ASGs that can subsequently be created on the new client's system. See the Reserved Indices topic (under Access Sub-Groups in this manual) for examples.

6. Click **OK** in the **Edit Access Sub-Group** dialog box to save your changes.

About Offline ASGs

Sometimes a client system may not be online when the HBAnyware Security Configurator is running. In this case, the Access Sub-Group (ASG) for the client appears offline in the ASG tree, much like the following:




Figure 52: The Access Sub-Groups Tab

The offline ASG entry serves as a placeholder for where the real ASG would be in the tree. You cannot modify or delete the entry (although it is removed from the display if all of its child ASGs are deleted).

It is possible to delete the child ASGs of an offline ASG. However, it is recommended that you delete them only if the client for the offline ASG will never be put online again. It is best to delete child ASGs when the parent ASG is online.

If you choose to delete a child ASG, the operation is immediate. There is no need to click **OK** or **Apply**.

Backup Masters

Introduction

A Backup Master mirrors the security data of the Master Security Client (MSC) in case it has to take over as the MSC if the MSC becomes unable to operate or is removed from the security configuration. A Backup master system receives all the updates to the security configuration on the MSC. However, you cannot make modifications to the security configuration on a Backup Master.

When the Configurator runs on a Backup Master, the **Access Control Group** tab looks like the tab on a non-MSC system. The **Access Sub-Group** tab displays the ASGs, but you cannot change the ASGs.

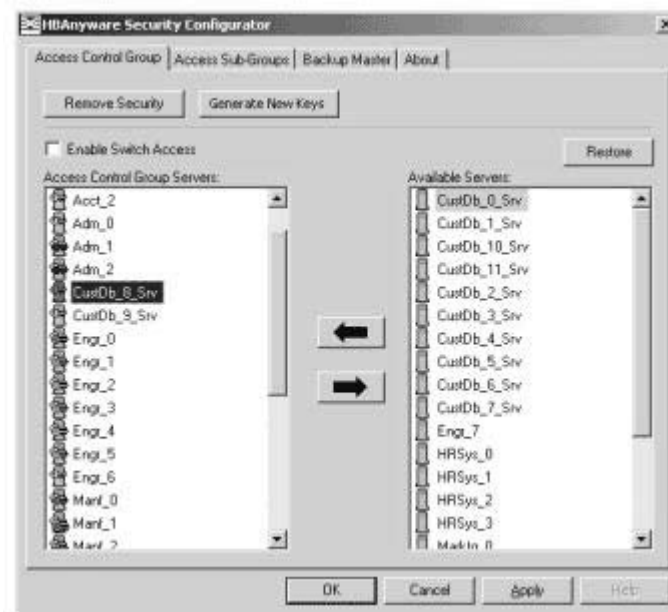


Figure 53: The Access Control Group Tab

The **Backup Master** tab is available only when the HBAAnyware Security Configurator is running on the MSC or a Backup Master. Use this tab to set up a system as a Backup Master to the MSC and to replace the MSC with a Backup Master.

Each time the HBAAnyware Security Configurator is started on the MSC and no Backup Master is assigned, a message warns you that no Backup Master Client is assigned to the security configuration.

If you run the HBAAnyware Security Configurator on a Backup Master, a message warns you that you can only view security information on a Backup Master. Security changes must be made to the MSC.

Because a Backup Master system receives all the updates that the MSC makes to the security configuration, it is very important that the Backup Master is online when the HBAAnyware Security Configurator is running on the MSC. Otherwise, updates to the security configuration are not reflected to the Backup Master. If the Backup Master then becomes the MSC, the security configuration may be corrupted.

Backup Master Eligible Systems

In order to be eligible to become a Backup Master, a system must not be a client or server in any ASG. In other words, it must be either a server in the MSC's Access Control Group (ACG) or an unsecure system. If it is an unsecure system, it will be secure when it becomes a Backup Master.

Backup Master Tab and Controls

The first time the **Backup Master** tab is selected on the MSC, it looks similar to the following:

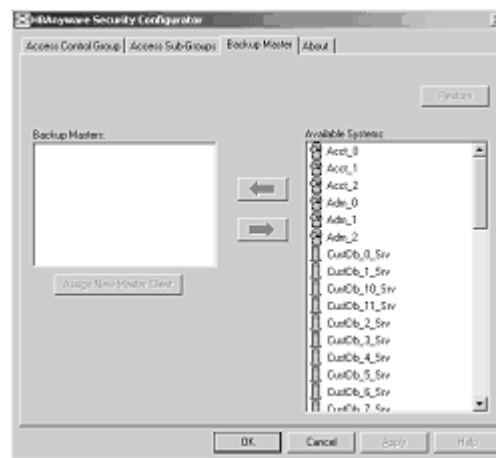


Figure 54: The Backup Master tab

Creating a Backup Master

To create a Backup Master:

1. On the Master Security Client (MSC), start the HBAAnyware Security Configurator.
2. Click the **Backup Master** tab.



Figure 55: The Backup Master Tab

3. Select a system from the Available Systems list.
4. Click the **left arrow** to move the system to the Backup Masters list.
5. Click **OK** or **Apply** to save your changes.

Reassigning a Backup Master as the New MSC from the Old MSC

Because a Backup Master may have to take over as the Master Security Client (MSC), it should be able to physically access all of the HBAs that the MSC can access. If the MSC is connected to multiple fabrics, its Backup Master should be selected from the Available Systems list that is connected to the same fabrics as the MSC.

To reassign a Backup Master as the new MSC from the old MSC:

1. On the MSC, start the HBAAnyware Security Configurator.
2. Click the **Backup Master** tab.

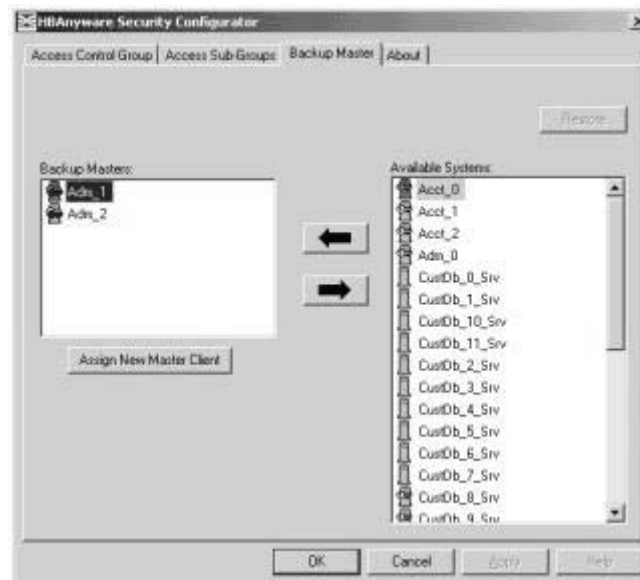


Figure 56: The Backup Master Tab

3. In the Backup Masters list, select the Backup Master system that you want to reassign as the MSC.
4. Click **Assign New Master Client**. You will be asked if you wish to proceed.
5. Click **Yes**. The selected Backup Master becomes the new MSC. The current MSC becomes a server in the new MSC's ACG. After the changes are made, a message indicates that the reassignment is complete.
6. Click **OK**. The Configurator closes because the system is no longer the MSC.

Reassigning a Backup Master as the New MSC from the Backup Master

WARNING: Use this method only if the MSC cannot relinquish control to a Backup Master. For example, if the MSC is no longer bootable or able to connect to the Fibre Channel network. Under any other circumstances, if the Backup Master takes over as the MSC, and the MSC is still running or comes back online later, there will be two MSCs for the same security configuration. This will eventually lead to corruption of the security configuration.

To reassign a Backup Master as the new MSC from the Backup Master:

1. On the Backup Master system that you want to reassign as the MSC, start the HBAAnyware Security Configurator.

2. Click the **Backup Master** tab.




Figure 57: The Backup Master “Warning” Dialog Box

3. Click **Assign This System As The Master Client**. A prompt asks if you want to continue.
4. Click **Yes**. A prompt notifies you that this system is now the new MSC.
5. Click **OK**. The Configurator closes. Restart the HBAAnyware Security Configurator to run the former Backup Master as the MSC.

Troubleshooting

Introduction

There are several circumstances in which your system may operate in an unexpected manner. The Troubleshooting section explains many of these circumstances and offers one or more workarounds for each situation.

Unusual Situations and Their Resolutions

General Situations

Table 3: General Situations

| Situation | Resolution |
|---|--|
| <p>Operating Error Occurs When Attempting to Run HBAnyware. When you attempt to run HBAnyware, an operating system error may occur. The computer may freeze.</p> | <p>Reboot the system.</p> |
| <p>Cannot See Multiple Zones from the Management Server. Cannot see multiple zones on the same screen of my management server running HBAnyware.</p> | <p>Provide a physical Fibre Channel connection into each of the zones. For each zone you want to see, connect an Emulex HBAnyware enabled port into that zone.</p> |
| <p>Cannot See Other HBAs or Hosts. Although HBAnyware is installed, only local host bus adapters (HBAs) are visible. The other HBAs and hosts in the SAN cannot be seen.</p> | <p>HBAnyware uses in-band data communication, meaning that the management server running HBAnyware must have a physical Fibre Channel connection to the SAN. All the HBAs in the SAN will be visible if:</p> <ul style="list-style-type: none"> • The other servers have a Fibre Channel connection to your zone of the SAN. Check fabric zoning. • Ensure that rmserver processes are running on the remote host: enter <code>ps -ef grep rmserver</code>). • All other HBAs are running HBAnyware and the appropriate driver. • The other HBAs are Emulex HBAs. <p>Note: HBAnyware must be running on all remote hosts that are to be discovered and managed. Remote capabilities of HBAnyware are subject to fabric zoning configuration. Remote hosts to be discovered and managed by HBAnyware must be in the same zone.</p> |
| <p>SAN Management Workstation Does Not Have a Fibre Channel Connection. The SAN management workstation does not have a physical Fibre Channel connection into the SAN because the other management tools are all out-of-band. Can HBAnyware be run on this SAN management workstation?</p> | <p>From the SAN management workstation, run a terminal emulation session into one of the servers that has HBAnyware loaded on it. Open an X-Windows session to run the server's HBAnyware GUI remotely.</p> |

Table 3: General Situations (Continued)

| Situation | Resolution |
|--|--|
| <p>Cannot See New LUNs. Although new LUNs were created on the storage array, they do not appear in HBAnyware.</p> | <ul style="list-style-type: none"> • Refresh the screen. |
| <p>HBAnyware Appears on Remote Servers in the SAN.</p> | <p>To prevent HBAnyware from appearing on remote servers in the SAN, do one of the following:</p> <p>Disable the rmserver process:</p> <ol style="list-style-type: none"> 1. Navigate to /usr/sbin/hbanyware. 2. Run ./stop_hbanyware to stop both the rmserver and elxdiscovery processes. 3. Run ./start_rmserver and ./start_elxdiscovery to restart both processes. <p>Disabling this service or process prevents the local servers from being seen remotely.</p> |
| <p>The HBAnyware Security Configurator (Security Configurator) software package will not install. An error message states that the latest version of HBAnyware must be installed first.</p> | <p>The system either has no HBAnyware software installed or has an older version of the HBAnyware software installed. In either case, obtain the latest version of the HBAnyware software and follow the installation instructions. Remember to install the HBAnyware software before installing the Security Configurator package.</p> |
| <p>Cannot access formerly accessible servers via the Security Configurator or the HBAnyware Utility.</p> | <p>This is actually a symptom of two different problems.</p> <ul style="list-style-type: none"> • New Keys Were Generated While Servers Were Offline • Security Removed While Servers Were Offline <p>See Table 8 on page 112 for details regarding these problems.</p> |
| <p>Cannot run the Security Configurator on a system that is configured for only secure access. I cannot run the Security Configurator on a system that is configured for only secure server access (it has no client privileges). The following message is displayed when the Security Configurator starts: "This system is not allowed client access to remote servers. This program will exit."</p> | <p>You cannot run the Security Configurator on a system that is configured for only secure server access. Click OK to close the message and the Configurator stops.</p> |

Security Configurator Situations - Access Control Groups (ACG)

Table 4: Access Control Groups Situations

| Situation | Resolution |
|--|---|
| <p>All servers are not displayed. When I run the Security Configurator on the Master Security Client (MSC), I do not see all of the systems in available servers or ACG Servers lists. When I run the Security Configurator on a non-MSC, I do not see all of the systems I should see in the ACG Servers list.</p> | <p>Make sure all of the systems are connected to the Fibre Channel network and are online when you start the Configurator. Discovery of the systems is done only once, at startup. Unlike the HBAnyware utility, there is no Rediscover Devices button. Therefore, the Security Configurator must be restarted to rediscover new systems.</p> |
| <p>Cannot add or remove a server. The Security Configurator shows only a list of the systems in this system's ACG. I cannot add or remove systems from the ACG.</p> | <p>This is normal. You can modify the ACG for your system only on the MSC or on a parent client system.</p> |
| <p>HBAnyware Utility shows non-ACG Servers. The HBAnyware utility shows servers that are part of the ACG and that are not part of the ACG.</p> | <p>The HBAnyware Utility discovers unsecured servers as well as servers that are part of its ACG. The servers that you see that are not part of the ACG are unsecured. They will be discovered by any system running the HBAnyware Utility on the same Fibre Channel fabric.</p> |

Security Configuration Situations - Access Sub-Groups (ASG)

Table 5: HBAware Security Configurator - Access Sub-Groups Situations

| Situation | Resolution |
|--|--|
| <p>ASG Appears to Be Non-Hierarchical. It is possible from a higher-level client (such as the MSC) to create an ASG 1 with system A as the client and systems B, C, D, and E as servers. Then create an ASG 2 with system E as the client, but with systems F and G as servers even though F and G are not part of ASG 1. This makes the topology non-hierarchical.</p> | <p>See “Non-Hierarchical and Hierarchical ASG” on page 113 for a discussion and a resolution to this situation.</p> |
| <p>Cannot add or remove a server.</p> | <p>When all of the systems in an ACG are running on a single fabric, they are all available to be added to any ASG. However, if the client is connected to more than one fabric, it is possible that not all of the servers in the client's ACG are physically accessible by a chosen client for an ASG. In this case, those servers are not available to be added to that ASG.</p> <p>If you add a system to an ASG as a server, and then make the system a client to a child ASG, you cannot remove it from the ACG it belongs to as a server until you delete the ASG to which it is a client.</p> <p>Before you delete a server from an ASG, you must first remove the server from any lower level ASGs to which it belongs.</p> |
| <p>In the ASG tree of the Access Sub-Groups tab, one or more of the names of the ASGs is displayed as "- ASG (Client Offline) -".</p> | <p>The client system for the ASG was not discovered when the Configurator was started. This is actually a symptom of two different problems.</p> <ul style="list-style-type: none"> • All Servers Are Not Displayed • New Keys Were Generated While Servers Were Offline <p>See Table 8 on page 112 for details regarding these problems.</p> |
| <p>Not All Servers are available to an ASG. When you create a new ASG or modify an existing ASG, not all of the servers in the ACG are available to be added to the ASG.</p> | <p>A client system can be connected to more than one fabric. While the system the Security Configurator is running on may be able to access all of the servers in its ACG, it is not necessarily the case that the selected client for the ASG can access all of the servers. Only those that can be accessed by the selected server will be available.</p> |

HBAnyware Security Configurator Situations - Backup Masters

Table 6: HBAnyware Security Configurator - Backup Masters Situations

| Situation | Resolution |
|--|--|
| <p>Cannot create a backup master.</p> | <p>Select a system (or group of systems) from the MSC to be the Backup Master. The system must be either an unsecured system (which will be secured by being made a Backup Master), or a system that is not part of any ASG (client or server). These systems will mirror the MSC's security configuration.</p> <p>Because the Backup Master may some day take over as the MSC, the Backup Master must be able to physically access all of the systems that the MSC can access. Therefore, if the MSC is connected to multiple fabrics, the Backup Master also must be connected to those same fabrics. When you select a Backup Master, the HBAnyware Security Configurator displays a warning if it detects that the system selected to be a Backup Master is not able to physically access the same systems that the MSC can access</p> |
| <p>Cannot modify the Security Configurator.</p> | <p>Select a system (or group of systems) from the MSC to be the Backup Master. The system must be either an unsecured system (which will be secured by being made a Backup Master), or a system that is not part of any ASG (client or server). These systems will mirror the MSC's security configuration.</p> <p>The Backup Master has client access from the HBAnyware Utility to all of the servers in the MSC's ACG. However, the Backup Master does not have client access to the MSC and it cannot modify the security configuration (create, modify or delete ASGs).</p> |
| <p>No Backup Master and the MSC Is no longer available. I do not have a Backup Master and the MSC system is no longer available. The servers are still secure. I installed the Security Configurator on another system, but I cannot access those servers to remove the security from them.</p> | <p>The servers are no longer part of a valid security configuration because there is no MSC to provide master control of the configuration. In order to reset the security on the affected servers, you will need to contact Emulex Technical Support to receive a special application and instructions on the reset procedure. After the servers have been reset, they should be seen by the Security Configurator and the HBAnyware Utility. At this point, you can set up security again through another MSC. At this time, also create a Backup Master.</p> |
| <p>The Backup Master tab is not available.</p> | <p>The Backup Master tab is displayed only when the Security Configurator is running on the MSC or a Backup Master. You use this tab to set up a system or systems to be backups to the MSC and to replace the MSC with a Backup Master.</p> <p>Each time you start the Security Configurator on the MSC and there is no Backup Master assigned, a warning message urges you to assign at least one Backup Master to prevent the loss of security information if the MSC were to become disabled.</p> |

Error Message Situations

Table 7: Error Message Situations

| Situation | Resolution |
|---|--|
| <p>The following error message is displayed when creating an ASG: "The Access Sub-Group name already exists. Please use a different name."</p> | <p>You entered a duplicate ASG name in the Access Sub-Group Name field. At each level of the security topology, each ASG name must be unique. Click OK on the message and enter a unique ASG name.</p> |
| <p>The following error message is displayed when deleting an ASG: "The Access Sub-Group parent's ASG is offline. You should delete the ASG when the parent ASG is available. This ASG should only be deleted if the parent ASG will not be available again. Are you sure you want to delete this ASG?"</p> | <p>The offline ASG entry serves as a placeholder for where the real ASG would be in the tree. You can neither modify nor delete it (although it is removed from the display if all of the child ASGs are deleted). It is possible to delete the child ASGs of the offline ASG. However, it is recommended that you delete them only if the client for the offline ASG will never be put online again. It is best to delete child ASGs when the parent ASG is online. Click Yes on the error message to delete the ASG or No to close the message without deleting.</p> |
| <p>The following error message is displayed when starting the HBAnyware Security Configurator: "This system is not allowed client access to remote servers. This program will exit."</p> | <p>The system you are running the Security Configurator on is already under the security umbrella as a server to one or more clients. To make this server a client (so that it can successfully run the Security Configurator), click OK to close the message and exit the program, then do the following:</p> <ol style="list-style-type: none"> 1. Run the Security Configurator on the MSC or on any client that has this server in its ASG. 2. Make this server a client to a group of servers. |
| <p>The following error message is displayed when starting the Security Configurator: "There are no Backup Master Client Systems assigned to this security configuration. At least one should be assigned to avoid loss of the security configuration should the Master Client System become disabled."</p> | <p>Use the Backup Master tab to assign a Backup Master for the MSC.</p> |
| <p>The first time the Security Configurator is started in an unsecure environment, the following message is displayed: "This utility is running on an unsecure system. Continuing will allow you to set up a new security configuration making this system the Master Client System."</p> | <p>Click OK on the message and complete the ACG setup. The system on which the Security Configurator is running will become the MSC.</p> |
| <p>When I start the Security Configurator on a Backup Master system, the following message is displayed: "Warning: This system is a backup master client system. Therefore you will only be able to view the security configuration. To make changes, you will need to run this utility on the master client system."</p> | <p>Because each Backup Master system receives all the updates that the MSC makes to the security configuration, the Backup Master systems must be online when the Security Configurator is running on the MSC. Otherwise, updates to the security configuration are not reflected to the Backup Master. If the Backup Master becomes the MSC, corruption of the security configuration may occur. Click OK to close the message.</p> |

Master Security Client Situations

Table 8: Master Security Client Situations

| Situation | Resolution |
|---|--|
| <p>The MSC is no longer bootable or able to connect to the FC network.</p> | <p>You must reassign a Backup Master as the new MSC from the Backup Master.</p> <p>Warning: Use this procedure only if the MSC cannot relinquish control to a Backup Master. For example, if the MSC is no longer bootable or able to connect to the FC network. Under any other circumstances, if the Backup Master takes over as the MSC and the MSC is still running or comes back online later, there will be two MSCs for the same security configuration. This will eventually lead to corruption of the security configuration.</p> |
| <p>New Keys Were Generated While Servers Were Offline. A "Generate New Keys" operation was performed while one or more of the servers were offline. Now those servers can no longer access the HBAnyware Security Configurator or the HBAnyware utility.</p> | <p>The servers are no longer part of the security configuration. In order to reset the security on the affected servers, you must contact Emulex Technical Support to receive a special application and instructions on the reset procedure. After the servers have been reset, they can be added back into the security topology by the MSC.</p> <p>Note: If the server was also a client to an ASG, then when you run the Security Configurator on the MSC or a parent client of this client, its label in the ASG tree of the Access Sub-Group tab will be "- ASG (Offline Client) -". You must delete the ASG (after deleting the child ASGs) and recreate the ASG configuration of this client and its child ASGs.</p> |
| <p>Security Removed While Servers Were Offline. Security was removed while one or more servers were offline. I can no longer access those servers from the Security Configurator or the HBAnyware utility.</p> | <p>The servers are no longer part of the security configuration. In order to reset the security on the affected servers, contact Emulex Technical Support to receive a special application and instructions on the reset procedure. After the servers have been reset, they should be seen by the Security Configurator or the HBAnyware utility.</p> |

Non-Hierarchical and Hierarchical ASG

It is possible from a higher-level client (such as the MSC) to create an ASG 1 with system A as the client and systems B, C, D, and E as servers. Then create an ASG 2 with system E as the client, but with systems F and G as servers even though F and G are not part of ASG 1. This makes the topology non-hierarchical (see Figure 58).

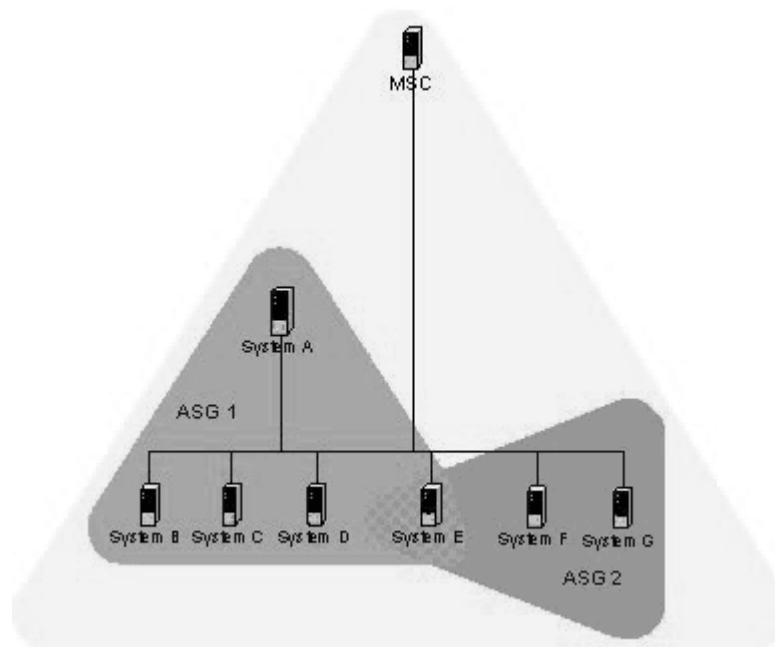


Figure 58: Non-hierarchical ASG Scenario

System E is part of ASG 1, but has been made a client of ASG 2, and both of the servers in ASG 2 are not part of ASG 1. You could not create this ASG on system A, but you could on the MSC (or on a parent client) because it can access systems F and G. Although not shown in the picture, it is also possible to make system A a server in ASG 2, creating a case where system A and system E are both clients and servers to/of each other.

While the Security Configurator will allow you to set up ASGs this way, it is best not to create a topology like this as it can lead to confusion. The best way is to set up the ASG on the MSC (or a higher-level parent) where the clients and servers do not cross over into other ASGs. Then set up ASGs on clients of those ASGs in the same manner, keeping the topology hierarchical (see Figure 59)




Figure 59: Hierarchical ASG Scenario

Ipfc Log Messages

Introduction

Log messages are organized into logical groups based on code functionality within the Fibre Channel driver. Each group consists of a block of 100 log message numbers. Most groups require a single block of 100 message numbers, however some groups (INIT, FCP) require two blocks.

The groups and the associated number ranges are defined in the Message Log table below. The preamble string shown in the Message Log table is displayed as part of the log message. The lower case 'x' of the preamble string defines the severity of the log message. The 'x' will be replaced by one of five lower case letters. Those letters are defined in the Severity Code table.

Severity Codes

Information and warning messages can be turned ON or OFF by setting/resetting the appropriate mask bit(s) in the variable 'log-verbose' located in the driver configuration module, lpfc.conf.c. By default, both information and warning messages are disabled. Configuration error (c), error (e), and panic (p) messages can not be disabled.

Table 9: Severity Code Table

| Code | Severity |
|------|---------------------|
| i | Information |
| w | Warning |
| c | Configuration Error |
| e | Error |
| p | Panic |

Message Group Masks

Table 10 defines the log message groups and the associated number ranges.

- The preamble string shown in this table is displayed as part of the log message.
- The lower case 'x' of the preamble string defines the severity of the log message and represents one of five lower case letters defined in the severity codes table.

Table 10: Message Log Table

| LOG Message Verbose Mask Definition | Preamble String | From | To | Verbose Bit | Verbose Description |
|-------------------------------------|-----------------|--------------|--------------|---------------------|-----------------------|
| ELS | ELx | 0100 | 0199 | LOG_ELS (0x1) | ELS events |
| DISCOVERY | Dlx | 0200 | 0299 | LOG_DISCOVERY (0x2) | Link discovery events |
| MBOX | MBx | 0300 | 0399 | LOG_MBOX (0x4) | Mailbox events |
| SLI | SLx | 0300 | 0399 | LOG_SLI (0x800) | SLI events |
| INIT | INx | 0400 0500 | 0499 0599 | LOG_INIT (0x8) | Initialization events |
| IP | IPx | 0600 | 0699 | LOG_IP (0x20) | IP traffic history |
| FCP | FPx | 0700 0800 | 0799 0899 | LOG_FCP (0x40) | FCP traffic history |
| NODE | NDx | 0900 | 0999 | LOG_NODE (0x80) | Node table events |
| Reserved | | 1000 | 1099 | Reserved | Reserved |
| Reserved | | 1100 | 1199 | Reserved | Reserved |

Table 10: Message Log Table (Continued)

| LOG Message Verbose Mask Definition | Preamble String | From | To | Verbose Bit | Verbose Description |
|-------------------------------------|-----------------|------|------|-----------------------|----------------------|
| MISC | Mlx | 1200 | 1299 | LOG_MISC (0x400) | Miscellaneous events |
| LINK | LKx | 1300 | 1399 | LOG_LINK_EVENT (0x10) | Link events |
| Reserved | | 1400 | 1499 | Reserved | Reserved |
| Reserved | | 1500 | 1599 | Reserved | Reserved |
| IOCTL | IOx | 1600 | 1699 | LOG_IOC (0x2000) | IOCTL events |
| All Messages | | | | LOG_ALL_MSG (0x1fff) | Log all messages |

Message Log Example

The following is an example of a LOG message:

```
Jul 12 16:30:26 <node> kernel: !lpfc0:0234:DIi:Device Discovery
completes
```

In the above LOG message:

- lpfc0 identifies the LOG message as coming from EMULEX HBA0.
- 0234 identifies the LOG message number.
- DIi identifies the LOG message as a DISCOVERY (DI) INFORMATION (i) message.

Note: If the word 'Data:' is present in a LOG message, any information to the right of 'Data:' is intended for Emulex technical support/engineering use only.

ELS Events (0100 - 0199)

0100 ELi: FLOGI failure

DESCRIPTION: An ELS FLOGI command that was sent to the fabric failed.

DATA: (1) ulpStatus (2) ulpWord[4]

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0101 ELi: FLOGI completes successfully

DESCRIPTION: An ELS FLOGI command that was sent to the fabric succeeded.

DATA: (1) ulpWord[4] (2) e_d_tov (3) r_a_tov (4) edtovResolution

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0102 ELi: PLOGI completes to NPort <nlp_DID>

DESCRIPTION: The HBA performed a PLOGI into a remote NPort.

DATA: (1) ulpStatus (2) ulpWord[4] (3) disc (4) num_disc_nodes

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0103 ELi: PRLI completes to NPort <nlp_DID>

DESCRIPTION: The HBA performed a PRLI into a remote NPort.

DATA: ((1) ulpStatus (2) ulpWord[4] (3) num_disc_nodes

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0104 ELi: ADISC completes to NPort <nlp_DID>

DESCRIPTION: The HBA performed a ADISC into a remote NPort.

DATA: (1) ulpStatus (2) ulpWord[4] (3) disc (4) num_disc_nodes

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0105 ELi: LOGO completes to NPort <nlp_DID>

DESCRIPTION: The HBA performed a LOGO to a remote NPort.

DATA: (1) ulpStatus (2) ulpWord[4] (3) num_disc_nodes

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0106 ELi: ELS cmd tag <ulploTag> completes

DESCRIPTION: The specific ELS command was completed by the firmware.

DATA: (1) ulpStatus (2) ulpWord[4]

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0107 ELi: Retry ELS command <elsCmd> to remote NPORT <did>

DESCRIPTION: The driver is retrying the specific ELS command.

DATA: ((1) retry (2) delay

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0108 ELi: No retry ELS command <elsCmd> to remote NPORT <did>

DESCRIPTION: The driver decided not to retry the specific ELS command that failed.

DATA: (1) retry (2) nlp_flag

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0109 ELi: ACC to LOGO completes to NPort <nlp_DID>

DESCRIPTION: The driver received a LOGO from a remote NPort and successfully issued an ACC response.

DATA: (1) nlp_flag (2) nlp_state (3) nlp_rpi

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0110 ELi: ELS response tag <ulploTag> completes

DESCRIPTION: The specific ELS response was completed by the firmware.

DATA: (1) ulpStatus (2) ulpWord[4] (3) nlp_DID (4) nlp_flag (5) nlp_state (6) nle.nlp_rpi

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0111 ELe: Dropping received ELS cmd

DESCRIPTION: The driver decided to drop an ELS Response ring entry.

DATA: (1) ulpStatus (2) ulpWord[4]

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver or firmware problem. If problems persist report these errors to Technical Support.

0112 ELi: ELS command <elsCmd> received from NPORT <did>

DESCRIPTION: Received the specific ELS command from a remote NPort.

DATA: (1) fc_ffstate

SEVERITY: Information

LOG: LOG_ELS Verbose

MODULE: fcelsb.c

ACTION: No action needed, informational.

0113 ELe: An FLOGI ELS command <elsCmd> was received from DID <did> in Loop Mode

DESCRIPTION: While in Loop Mode an unknown or unsupported ELS command was received.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: Check device DID

0114 ELi: FLOGI/PLOGI chkparm OK

DESCRIPTION: Received a FLOGI/PLOGI from a remote NPORT and its Fibre Channel service parameters match this HBA. Request can be accepted.

DATA: (1) nlp_DID (2) nlp_state (3) nlp_flag (4) nlp_Rpi

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0115 ELe: Unknown ELS command <elsCmd> received from NPORT <did>

DESCRIPTION: Received an unsupported ELS command from a remote NPORT.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: Check remote NPORT for potential problem.

0116 ELi: Xmit ELS command <elsCmd> to remote NPORT <did>

DESCRIPTION: Xmit ELS command to remote NPORT.

DATA: (1) icmd->ulploTag (2) binfo->fc_ffstate

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0117 ELi: Xmit ELS response <elsCmd> to remote NPORT <did>

DESCRIPTION: Xmit ELS response to remote NPORT.

DATA: (1) icmd->ulploTag (2) size

SEVERITY: Information

LOG: LOG_ELS Verbose

MODULE: fcelsb.c

ACTION: No action needed, informational.

0118 ELi: Xmit CT response on exchange <xid>

DESCRIPTION: Xmit a CT response on the appropriate exchange.

DATA: (1) ulploTag (2) fc_ffstate

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0119 ELi: Issue GEN REQ IOCB for NPORT <did>

DESCRIPTION: Issue a GEN REQ IOCB for remote NPORT. These are typically used for CT request.

DATA: (1) ulploTag (2) fc_ffstate

SEVERITY: Information

LOG: LOG_ELS Verbose

ACTION: No action needed, informational.

0127 ELe: ELS timeout

DESCRIPTION: An ELS IOCB command was posted to a ring and did not complete within ULP timeout seconds.

DATA: (1) elscmd (2) did (3) ulpcommand (4) iotag

SEVERITY: Error

LOG: Always

ACTION: If no ELS command is going through the adapter, reboot the system; If problem persists, contact Technical Support.

Link Discovery Events (0200 - 0299)

0200 Dle: CONFIG_LINK bad hba state <hba_state>

DESCRIPTION: A CONFIG_LINK mbox command completed and the driver was not in the right state.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: Software driver error. If this problem persists, report these errors to Technical Support.

0201 Dli: Abort outstanding I/O on NPort <nlp_DID>

DESCRIPTION: All outstanding I/Os are cleaned up on the specified remote NPort.

DATA: (1) nlp_flag (2) nlp_state (3) nle.nlp_rpi

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0202 Dli: Start Discovery hba state <hba_state>

DESCRIPTION: Device discovery / rediscovery after FLOGI, FAN or RSCN has started.

DATA: (1) tmo (2) fc_plogi_cnt (3) fc_adisc_cnt

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0203 Dle: Nodev timeout on NPort <nlp_DID>

DESCRIPTION: A remote NPort that was discovered by the driver disappeared for more than ELX_NODEV_TMO seconds.

DATA: (1) nlp_flag (2) nlp_state (3) nlp_rpi

SEVERITY: Error

LOG: Always

ACTION: Check connections to Fabric / HUB or remote device.

0204 Dli: Create SCSI Target <tgt>

DESCRIPTION: A mapped FCP target was discovered and the driver has allocated resources for it.

DATA: None

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: No action needed, informational.

0205 Dli: Create SCSI LUN <lun> on Target <tgt>

DESCRIPTION: A LUN on a mapped FCP target was discovered and the driver has allocated resources for it.

DATA: None

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: No action needed, informational.

0206 Dli: Report Lun completes on NPort <nlp_DID>

DESCRIPTION: The driver issued a REPORT_LUN SCSI command to a FCP target and it completed.

DATA: (1) ulpStatus (2) rspStatus2 (3) rspStatus3 (4) nlp_failMask

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: No action needed, informational.

0207 Dli: Issue Report LUN on NPort <nlp_DID>

DESCRIPTION: The driver issued a REPORT_LUN SCSI command to a FCP target.

DATA: (1) nlp_failMask (2) nlp_state (3) nlp_rpi

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

No action needed, informational.

0208 Dli: Failmask change on NPort <nlp_DID>

DESCRIPTION: An event was processed that indicates the driver may not be able to communicate with the remote NPort.

DATA: (1) nlp_failMask (2) bitmask (3) flag

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0209 Dli: RFT request completes ulpStatus <ulpStatus> CmdRsp <CmdRsp>

DESCRIPTION: A RFT request that was sent to the fabric completed.

DATA: (1) nlp_failMask (2) bitmask (3) flag

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0210 Dli: Continue discovery with <num_disc_nodes> ADISCS to go

DESCRIPTION: A device discovery is in progress.

DATA: (1) fc_adisc_cnt (2) fc_flag (3) phba->hba_state

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0211 Dli: DSM in event <evt> on NPort <nlp_DID> in state <cur_state>

DESCRIPTION: The driver Discovery State Machine is processing an event.

DATA: (1) nlp_flag

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0212 Dli: DSM out state <rc> on NPort <nlp_DID>

DESCRIPTION: The driver Discovery State Machine completed processing an event.

DATA: (1) nlp_flag

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0213 Dli: Reassign scsi id <sid> to NPort <nlp_DID>

DESCRIPTION: A previously bound FCP Target has been rediscovered and reassigned a scsi id.

DATA: (1) nlp_bind_type (2) nlp_flag (3) nlp_state (4) nlp_rpi

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: No action needed, informational.

0214 Dli: RSCN received

DESCRIPTION: An RSCN ELS command was received from a fabric.

DATA: (1) fc_flag (2) i (3) *lp (4) fc_rscn_id_cnt

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0215 Dli: RSCN processed

DESCRIPTION: An RSCN ELS command was received from a fabric and processed.

DATA: (1) fc_flag (2) cnt (3) fc_rscn_id_cnt (4) fc_ffstate

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0216 Dli: Assign scandown scsi id <sid> to NPort <nlp_DID>

DESCRIPTION: A scsi id is assigned due to BIND_ALPA.

DATA: ((1) nlp_bind_type (2) nlp_flag (3) nlp_state (4) nlp_rpi

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: No action needed, informational.

0217 Dli: Unknown Identifier in RSCN payload

DESCRIPTION: Typically the identifier in the RSCN payload specifies a domain, area or a specific NportID. If neither of these are specified, a warning will be recorded.

DATA: (1) didp->un.word

SEVERITY: Error

LOG: Always

ACTION: Potential problem with Fabric. Check with Fabric vendor.

0218 Dli: FDMI Request

DESCRIPTION: The driver is sending an FDMI request to the fabric.

DATA: (1) fc_flag (2) hba_state (3) cmdcode

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0219 Dli: Issue FDMI request failed

DESCRIPTION: Cannot issue FDMI request to HBA.

DATA: (1) cmdcode

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0220 Dli: FDMI rsp failed

DESCRIPTION: An error response was received to FDMI request.
DATA:(1) SWAP_DATA16 (fdmi_cmd)
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: The fabric does not support FDMI, check fabric configuration.

0221 DIw: FAN timeout

DESCRIPTION: A link up event was received without the login bit set, so the driver waits E_D_TOV for the Fabric to send a FAN. If no FAN is received, a FLOGI will be sent after the timeout.
DATA: None
SEVERITY: Warning
LOG: LOG_DISCOVERY Verbose
ACTION: None required. The driver recovers from this condition by issuing a FLOGI to the fabric.

0222 DIe: Initial FLOGI timeout

DESCRIPTION: The driver sent the initial FLOGI to fabric and never got a response back.
DATA: None
SEVERITY: Error
LOG: Always
ACTION: Check Fabric configuration. The driver recovers from this and continues with device discovery.

0223 DIe: Timeout while waiting for NameServer login

DESCRIPTION: Our login request to the NameServer was not acknowledged within RATOV.
DATA: None
SEVERITY: Error
LOG: Always
ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

0224 DIe: NameServer Query timeout

DESCRIPTION: Node authentication timeout, node Discovery timeout. A NameServer Query to the Fabric or discovery of reported remote NPorts is not acknowledged within R_A_TOV.
DATA: (1) fc_ns_retry (2) fc_max_ns_retry
SEVERITY: Error
LOG: Always
ACTION: Check Fabric configuration. The driver recovers from this and continues with device discovery.

0225 Dli: Device Discovery completes

DESCRIPTION: This indicates successful completion of device (re)discovery after a link up.
DATA: None
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0226 Dle: Device discovery completion error

DESCRIPTION: This indicates that an uncorrectable error was encountered during device (re)discovery after a link up. Fibre Channel devices will not be accessible if this message is displayed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: Reboot the system. If the problem persists, report the error to Technical Support. Run with Verbose mode on for more details.

0227 Dle: Node Authentication timeout

DESCRIPTION: The driver has lost track of what NPORTs are being authenticated.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: None required. The driver should recover from this event.

0228 Dle: CLEAR LA timeout

DESCRIPTION: The driver issued a CLEAR_LA that never completed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: None required. The driver should recover from this event.

0229 Dli: Assign scsi ID <sid> to NPort <nlp_DID>

DESCRIPTION: The driver assigned a scsi id to a discovered mapped FCP target.

DATA: (1) nlp_bind_type (2) nlp_flag (3) nlp_state (4) nlp_rpi

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: No action needed, informational.

0230 Dli: Cannot assign scsi ID on NPort <nlp_DID>

DESCRIPTION: The driver cannot assign a scsi id to a discovered mapped FCP target.

DATA: (1) nlp_flag (2) nlp_state (3) nlp_rpi

SEVERITY: Information

LOG: LOG_DISCOVERY | LOG_FCP Verbose

ACTION: Check persistent binding information.

0231 Dle: RSCN timeout

DESCRIPTION: The driver has lost track of what NPORTs have RSCNs pending.

DATA: (1) fc_ns_retry (2) fc_max_ns_retry

SEVERITY: Error

LOG: Always

ACTION: None required. The driver should recover from this event.

0232 Dli: Continue discovery with <num_disc_nodes> PLOGIs to go

DESCRIPTION: Device discovery is in progress.
DATA: (1) fc_plogi_cnt (2) fc_flag (3) phba->hba_state
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0234 Dli: ReDiscovery RSCN

DESCRIPTION: The number / type of RSCNs has forced the driver to go to the nameserver and re-discover all NPORTs.
DATA: (1) fc_defer_rscn.q_cnt (2) fc_flag (3) hba_state
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0235 Dli: Deferred RSCN

DESCRIPTION: The driver has received multiple RSCNs and has deferred the processing of the most recent RSCN.
DATA: (1) fc_defer_rscn.q_cnt (2) fc_flag (3) hba_state
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0236 Dli: NameServer req

DESCRIPTION: The driver is issuing a NameServer request to the fabric.
DATA: (1) cmdcode (2) fc_flag (3) fc_rscn_id_cnt
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0237 Dlw: Pending Link Event during Discovery

DESCRIPTION: Received link event during discovery. Causes discovery restart.
DATA: (1) hba_state
SEVERITY: Warning
LOG: LOG_DISCOVERY Verbose
ACTION: None required unless problem persists. If problem persists, check cabling.

0238 Dli: NameServer Rsp

DESCRIPTION: The driver received a NameServer response.
DATA: (1) Did (2) nlp_flag (3) fc_flag (4) fc_rscn_id_cnt
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0240 Dli: NameServer Rsp Error

DESCRIPTION: The driver received a NameServer response containing a status error.
DATA: (1) CommandResponse.bits.CmdRsp (2) ReasonCode (3) Explanation (4) fc_flag
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

0241 Dli: NameServer rsp error

DESCRIPTION: The driver received a NameServer response containing a status error.
DATA: (1) CommandResponse.bits.CmdRsp (2) ReasonCode (3) Explanation (4) fc_flag
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: Check the fabric configuration. The driver recovers from this and continues with device discovery.

0242 Dli: Failmask change on TGT <target_ID> LUN <lun_ID>

DESCRIPTION: An event was processed that indicates the driver may not be able to send I/O to the specified LUN.
DATA: (1) nlp_failMask (2) bitmask (3) flag
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, information.

0243 Dli: Issue FDMI request failed

DESCRIPTION: Cannot issue an FDMI request to HBA.
DATA: (1) cmdcode
SEVERITY: Information
LOG: LOG_DISCOVERY Verbose
ACTION: No action needed, informational.

0244 Dli: Issue FDMI request failed

DESCRIPTION: Cannot issue an FDMI request to the HBA.
DATA: (1) cmdcode
SEVERITY: Information
LOG: LOG_Discovery Verbose
ACTION: No action needed, informational.

0245 Dlw: ALPA based bind method used on an HBA which is in a nonloop topology

DESCRIPTION: ALPA-based bind method used on an HBA which is not in a loop topology.
DATA: (1) topology
SEVERITY: Warning
LOG: LOG_DISCOVERY Verbose
ACTION: Change the bind method configuration parameter of the HBA to 1(WWNN) or 2(WWPN) or 3(DID)

0246 Dle: RegLogin failed

DESCRIPTION: The firmware returned a failure for the specified RegLogin.

DATA: Did, mbxStatus, hbaState

SEVERITY: Error

LOG: Always

MODULE: fcscsib.c

ACTION: This message indicates that the firmware could not do RegLogin for the specified Did. There may be a limitation on how many nodes an HBA can see.

0247 Dli: Start Discovery Timer state <hba_state>

DESCRIPTION: Start the device discovery / RSCN rescue timer.

DATA: (1) tmo (2) disctmo (3) fc_plogi_cnt (4) fc_adisc_cnt

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0248 Dli: Cancel Discovery Timer state <hba_state>

DESCRIPTION: Cancel the device discovery / RSCN rescue timer.

DATA: (1) fc_flag (2) rc (3) fc_plogi_cnt (4) fc_adisc_cnt

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

0249 Dli: Start nodev Timer

DESCRIPTION: A device disappeared from the FC network. If the device does not return within the nodev-tmo timeout, all I/O to the device will fail.

DATA: (1) nlp_DID (2) nlp_flag (3) nlp_state (4) nlp

SEVERITY: Information

LOG: LOG_DISCOVERY Verbose

ACTION: No action needed, informational.

Mailbox Events (0300 - 0399)

0300 MBw: READ_LA: no buffers

DESCRIPTION: The driver attempted to issue a READ_LA mailbox command to the HBA, but there were no buffers available.

DATA: None

SEVERITY: Warning

LOG: LOG_MBOX Verbose

ACTION: This message indicates: (1) a possible lack of memory resources. Try increasing the lpfc 'num_bufs' configuration parameter to allocate more buffers. (2) A possible driver buffer management problem. If this problem persists, report the error to Technical Support.

0301 MBw: READ_SPARAM: no buffers

DESCRIPTION: The driver attempted to issue a READ_SPARAM mailbox command to the HBA, but there were no buffers available.

DATA: None

SEVERITY: Warning

LOG: LOG_MBOX Verbose

ACTION: This message indicates: (1) a possible lack of memory resources. Try increasing the lpfc 'num_bufs' configuration parameter to allocate more buffers. (2) A possible driver buffer management problem. If the problem persists, report the error to Technical Support.

0302 MBw: REG_LOGIN: no buffers

DESCRIPTION: The driver attempted to issue a REG_LOGIN mailbox command to the HBA, but there were no buffers available.

DATA: None

SEVERITY: Warning

LOG: LOG_MBOX Verbose

ACTION: This message indicates: (1) a possible lack of memory resources. Try increasing the lpfc 'num_bufs' configuration parameter to allocate more buffers. (2) A possible driver buffer management problem. If the problem persists, report the error to Technical Support.

0303 INe: Adapter initialization error, mbxCmd <cmd> READ_NVPARM, mbxStatus <status>

DESCRIPTION: A mailbox command failed during initialization.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0304 MBe: Stray mailbox interrupt, mbxCommand <cmd> mbxStatus <status>

DESCRIPTION: Received a mailbox completion interrupt and there are no outstanding mailbox commands.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0305 MBi: Mbox cmd cmpl error - RETRYing

DESCRIPTION: A mailbox command completed with an error status that causes the driver to reissue the mailbox command.

DATA: (1) mbxCommand (2) mbxStatus (3) word1 (4) hba_state

SEVERITY: Information

LOG: LOG_MBOX Verbose

ACTION: No action needed, informational.

0306 MBe: CONFIG_LINK mbxStatus error <mbxStatus> HBA state <hba_state>

DESCRIPTION: The driver issued a CONFIG_LINK mbox command to the HBA that failed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a firmware or hardware problem. Report these errors to Technical Support.

0307 MBi: Mailbox Cmpl, wd0 <pmbox> wd1 <varWord> wd2 <varWord> cmpl <mbox_cmpl)

DESCRIPTION: A mailbox command completed.

DATA: None

SEVERITY: Information

LOG: LOG_MBOX Verbose

ACTION: No action needed, informational.

0308 MBi: Mbox cmd issue - BUSY

DESCRIPTION: The driver attempted to issue a mailbox command while the mailbox was busy processing the previous command. The processing of the new command will be deferred until the mailbox becomes available.

DATA: (1) mbxCommand (2) hba_state (3) sli_flag (4) flag

SEVERITY: Information

LOG: LOG_MBOX Verbose

ACTION: No action needed, informational.

0309 MBi: Mailbox cmd <cmd> issue

DESCRIPTION: The driver is in the process of issuing a mailbox command.

DATA: (1) hba_state (2) sli_flag (3) flag

SEVERITY: Information

LOG: LOG_MBOX Verbose

ACTION: No action needed, informational.

0310 MBe: Mailbox command <cmd> timeout

DESCRIPTION: A mailbox command was posted to the adapter and did not complete within 30 seconds.

DATA: (1) hba_state (2) sli_flag (3) mbox_active

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver or firmware problem. If no I/O is going through the adapter, reboot the system. If the problem persists, report the error to Technical Support.

0311 MBi: Mailbox command <cmd> cannot issue

DESCRIPTION: The driver is in the wrong state to issue the specified command.

DATA: (1) hba_state (2) sli_flag (3) flag

SEVERITY: Information

LOG: LOG_MBOX Verbose

MODULE: fcscsib.c

ACTION: No action needed, informational.

0312 SLe: Ring <ringno> handler: portRspPut <portRspPut> is bigger then rsp ring <portRspMax>

DESCRIPTION: The port rsp ring put index is larger than the size of the rsp ring.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0313 SLw: Ring <ringno> handler: unexpected Rctl <Rctl> Type <Type> received

DESCRIPTION: The Rctl/Type of a received frame did not match any for the configured masks for the specified ring.

DATA: None

SEVERITY: Warning

LOG: Always

ACTION: This warning could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0314 SLe: Ring <ringno> issue: portCmdGet <portCmdGet> is bigger then cmd ring <portCmdMax>

DESCRIPTION: The port cmd ring get index is greater than the size of cmd ring.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0315 SLe: Ring <ringno> issue: portCmdGet <portCmdGet> is bigger then cmd ring <portCmdMax>

DESCRIPTION: The port cmd ring get index is greater than the size of cmd ring.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0316 SLe: Cmd ring <ringno> put: iotag <iotag> greater then configured max <fast_iotag> wd0 <icmd>

DESCRIPTION: The assigned I/O iotag is greater than the allowed maximum.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0317 SLe: Rsp ring <ringno> get: iotag <iotag> greater then configured max <fast_iotag> wd0 <irsp>

DESCRIPTION: The assigned I/O iotag is greater than the maximum allowed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0318 SLi: Outstanding I/O count for ring <ringno> is at max <fast_iotag>

DESCRIPTION: An I/O tag cannot be assigned because none are available. The maximum number of allowed I/Os are currently outstanding.

DATA: None

SEVERITY: Information

LOG: LOG_SLI Verbose

ACTION: This message indicates the adapter HBA I/O queue is full. Typically this happens when heavy I/O is running on a low-end (3 digit) adapter. We suggest you upgrade to a high-end adapter.

0319 MBe: The driver issued a READ_SPARAM mbox command to the HBA that failed.

DESCRIPTION: The driver issued a READ_SPARAM mbox command to the HBA that failed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a firmware or hardware problem. Report these errors to Technical Support.

0320 MBe: CLEAR_LA mbxStatus error <mbxStatus> hba state <hba_state>

DESCRIPTION: The driver issued a CLEAR_LA mbox command to the HBA that failed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a firmware or hardware problem. Report these errors to Technical Support.

0321 SLe: Unknown IOCB command

DESCRIPTION: Received an unknown IOCB command completion.

DATA: (1) ulpCommand (2) ulpStatus (3) ulploTag (4) ulpContext

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver or firmware problem. If these problems persist, report these errors to Technical Support.

0322 SLw: Ring <ringno> handler: unexpected completion IoTag <IoTag>

DESCRIPTION: The driver could not find a matching command for the completion received on the specified ring.

DATA: (1) ulpStatus (2) ulpWord[4] (3) ulpCommand (4) ulpContext

SEVERITY: Warning

LOG: LOG_SLI Verbose

ACTION: This warning could indicate a software driver or firmware problem. If the problem persists report these errors to Technical Support.

0323 MBe: Unknown Mailbox command <cmd> Cmpl

DESCRIPTION: A unknown mailbox command completed.

DATA: (1) Mailbox Command

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0324 MBe: Adapter initialization error, mbxCmd <cmd> READ_NVPARM, mbxStatus <status>

DESCRIPTION: A read nvparams mailbox command failed during port configuration.

DATA:(1) Mailbox Command (2) Mailbox Command Status

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver, firmware or hardware problem. Report these errors to Technical Support.

0325 SLw: Post buffer for ring <num> failed

DESCRIPTION: The driver cannot allocate a buffer to post to the ring. This usually indicates that the host system is out of buffers.

DATA:(1) missbufcnt

SEVERITY: Warning

LOG: LOG_SLI verbose

ACTION: Report these errors to Technical Support if the problem persists.

Initialization Events (0400 - 0499)

0405 INi: Service Level Interface (SLI) 2 selected

DESCRIPTION: A CONFIG_PORT (SLI2) mailbox command was issued.

DATA: None

SEVERITY: Information

LOG: LOG_INIT Verbose

ACTION: No action needed, informational.

0406 INw: Memory buffer pool is below low water mark

DESCRIPTION: A driver memory buffer pool is low on buffers.

DATA: (1) seg (2) fc_lowmem (3) low

SEVERITY: Warning

LOG: LOG_INIT Verbose

ACTION: None required. The driver will recover as buffers are returned to the pool.

0407 INe: Memory Buffer Pool is at upper limit.

DESCRIPTION: A memory buffer pool cannot add more buffers because it is at its himem value.

DATA: (1) seg (2) q_cnt (3) himem

SEVERITY: Error

LOG: Always

ACTION: None required. The driver will recover as buffers are returned to the pool.

0409 INe: Memory Buffer Pool is out of buffers

DESCRIPTION: A driver memory buffer pool is exhausted.

DATA: (1) seg (2) fc_free (3) fc_mbox.q_cnt (4) fc_memhi

SEVERITY: Error

LOG: Always

ACTION: Configure more resources for that buffer pool. If the problem persists, report the error to Technical Support.

0410 INe: Cannot find virtual addr for mapped buf on ring <num>

DESCRIPTION: The driver cannot find the specified buffer in its mapping table. Thus it cannot find the virtual address needed to access the data.

DATA: (1) first (2) q_first (3) q_last (4) q_cnt

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver or firmware problem. If the problem persists report these errors to Technical Support.

0411 INc: fcp_bind_method is 4 with Persistent binding - ignoring fcp_bind_method

DESCRIPTION: The configuration parameter for fcp_bind_method conflicts with Persistent binding parameter.

DATA: (1) a_current (2) fcp_mapping

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0412 INe: Scan-down is out of range - ignoring scan-down

DESCRIPTION: The configuration parameter for scan-down is out of range.

DATA: (1) clp[CFG_SCAN_DOWN].a_current (2) fcp_mapping

SEVERITY: Error

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0413 es0: Configuration parameter out of range, resetting to default value

DESCRIPTION: You are attempting to set a configuration parameter to a value not supported by the driver. Resetting the configuration parameter to the default value.

DATA: (1) a_string (2) a_low (3) a_hi (4) a_default

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0427 INc: Same node has multiple persistent WWPN binding definitions.

DESCRIPTION: You are attempting to define multiple persistent WWPN bindings to a single node. Only the first persistent binding is accepted and the rest are ignored.

DATA: (1) a_string

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0428 INc: Same node has multiple persistent WWNN binding definitions.

DESCRIPTION: You are attempting to define multiple persistent WWNN bindings to a single node. Only the first persistent binding is accepted and the rest are ignored.

DATA: (1) a_string

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0429 INc: Same node has multiple persistent DID binding definitions.

DESCRIPTION: You are attempting to define multiple persistent DID bindings to a single node. Only the first persistent binding is accepted and the rest are ignored.

DATA: (1) a_string

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0430 INc: WWPN binding entry <num>: syntax error code <code>

DESCRIPTION: A syntax error occurred while parsing WWPN binding configuration information.

DATA: None

Detail: Binding syntax error codes

0 FC_SYNTAX_OK

1 FC_SYNTAX_OK_BUT_NOT_THIS_BRD

2 FC_SYNTAX_ERR_ASC_CONVERT

3 FC_SYNTAX_ERR_EXP_COLON

4 FC_SYNTAX_ERR_EXP_LPFC

5 FC_SYNTAX_ERR_INV_LPFC_NUM

6 FC_SYNTAX_ERR_EXP_T

7 FC_SYNTAX_ERR_INV_TARGET_NUM

8 FC_SYNTAX_ERR_EXP_D

9 FC_SYNTAX_ERR_INV_DEVICE_NUM

10 FC_SYNTAX_ERR_INV_RRATIO_NUM

11 FC_SYNTAX_ERR_EXP_NULL_TERM

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0431 INc: WWNN binding entry <num>: syntax error code <code>

DESCRIPTION: A syntax error occurred while parsing WWNN binding configuration information.

DATA: None

Detail: Binding syntax error codes

0 FC_SYNTAX_OK

1 FC_SYNTAX_OK_BUT_NOT_THIS_BRD

2 FC_SYNTAX_ERR_ASC_CONVERT

3 FC_SYNTAX_ERR_EXP_COLON

4 FC_SYNTAX_ERR_EXP_LPFC

5 FC_SYNTAX_ERR_INV_LPFC_NUM

6 FC_SYNTAX_ERR_EXP_T

7 FC_SYNTAX_ERR_INV_TARGET_NUM

8 FC_SYNTAX_ERR_EXP_D

9 FC_SYNTAX_ERR_INV_DEVICE_NUM

10 FC_SYNTAX_ERR_INV_RRATIO_NUM

11 FC_SYNTAX_ERR_EXP_NULL_TERM

SEVERITY: Error config

LOG: always

ACTION: Make necessary changes to the lpfc configuration file.

0432 INc: WWPN binding entry: node table full

DESCRIPTION: More bindings entries were configured than the driver can handle.

DATA: None

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file so that fewer bindings are configured.

0433 INc: WWNN binding entry: node table full

DESCRIPTION: More bindings entries were configured than the driver can handle.

DATA: None

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file so that fewer bindings are configured.

0434 INc: DID binding entry <num>: syntax error code <code>

DESCRIPTION: A syntax error occurred while parsing DID binding configuration information.

DATA: None

Detail: Binding syntax error codes

0 FC_SYNTAX_OK

1 FC_SYNTAX_OK_BUT_NOT_THIS_BRD

2 FC_SYNTAX_ERR_ASC_CONVERT

3 FC_SYNTAX_ERR_EXP_COLON

4 FC_SYNTAX_ERR_EXP_LPFC

5 FC_SYNTAX_ERR_INV_LPFC_NUM

6 FC_SYNTAX_ERR_EXP_T

7 FC_SYNTAX_ERR_INV_TARGET_NUM

8 FC_SYNTAX_ERR_EXP_D

9 FC_SYNTAX_ERR_INV_DEVICE_NUM

10 FC_SYNTAX_ERR_INV_RRATIO_NUM

11 FC_SYNTAX_ERR_EXP_NULL_TERM

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

0435 INc: DID binding entry: node table full

DESCRIPTION: More bindings entries were configured than the driver can handle.

DATA: None

SEVERITY: Error config

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file so that fewer bindings are configured.

0436: Adapter failed to init, timeout, status reg <status>

DESCRIPTION: The adapter failed during powerup diagnostics after it was reset.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0437 INe: Adapter failed to initialize chipset

DESCRIPTION: The adapter failed during powerup diagnostics after it was reset.

DATA: (1) status (2) status1 (3) status2

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0438 INe: Adapter failed to initialize chipset

DESCRIPTION: The adapter failed during powerup diagnostics after it was reset.

DATA: (1) status (2) status1 (3) status2

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0439 INe: Adapter failed to init, mbxCmd <cmd> READ_REV, mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing a READ_REV mailbox command.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0440 INe: Adapter failed to init, mbxCmd <cmd> READ_REV, detected outdated firmware

DESCRIPTION: Outdated firmware was detected during initialization.

DATA: (1) read_rev_reset

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. Update the firmware. If the problem persists, report the error to Technical Support.

0441 INi: VPD not present on adapter, mbxCmd <cmd> DUMP_VPD, mbxStatus <status>

DESCRIPTION: The DUMP_VPD mailbox command failed.

DATA: None

SEVERITY: Information

LOG: LOG_INIT Verbose

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0442 INe: Adapter failed to init, mbxCmd <cmd> CONFIG_PORT, mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing a CONFIG_PORT mailbox command.

DATA: (1) hbainit

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0446 INe: Adapter failed to init, mbxCmd <cmd> CFG_RING, mbxStatus <status>, ring <num>

DESCRIPTION: Adapter initialization failed when issuing a CFG_RING mailbox command.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0447 INe: Adapter failed init, mbxCmd <cmd> CONFIG_LINK mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing a CONFIG_LINK mailbox command.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0448 INe: Adapter failed to init, mbxCmd <cmd> READ_SPARM, mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing a READ_SPARM mailbox command.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0449 INe: WorldWide PortName type <type> doesn't conform to IP Profile

DESCRIPTION: In order to run IP, the WorldWide PortName must be of type IEEE (NAA = 1). This message is displayed if the adapter WWPN doesn't conform with the standard.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: Turn off the network-on configuration parameter or configure a different WWPN.

0450 INw: Adapter failed to init, mbxCmd <cmd> FARP, mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing a FARP mailbox command.

DATA: None

SEVERITY: Warning

LOG: LOG_INIT Verbose

ACTION: None required.

0451 INe: Enable interrupt handler failed

DESCRIPTION: The driver attempted to register the HBA interrupt service routine with the host operating system, but failed.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or driver problem. If the problem persists, report the error to Technical Support.

0453 INe: Adapter failed to init, mbxCmd <cmd> READ_CONFIG, mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing a READ_CONFIG mailbox command.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0454 INe: Adapter failed to init, mbxCmd <cmd> INIT_LINK, mbxStatus <status>

DESCRIPTION: Adapter initialization failed when issuing an INIT_LINK mailbox command.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0455 INi: Vital Product

DESCRIPTION: Vital product data (VPD) contained in the HBA flash.

DATA: (1) vpd[0] (2) vpd[1] (3) vpd[2] (4) vpd[3]

SEVERITY: Information

LOG: LOG_INIT Verbose

ACTION: No action needed, informational.

0457 INe: Adapter Hardware Error

DESCRIPTION: The driver received an interrupt indicting a possible hardware problem.

Data: (1) status (2) status1 (3) status2

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a hardware or firmware problem. If the problem persists, report the error to Technical Support.

0458 INw: Bring adapter online

DESCRIPTION: The FC driver has received a request to bring the adapter online. This may occur when running lputil.

DATA: None

SEVERITY: Warning

LOG: LOG_INIT Verbose

ACTION: None required.

0460 INw: Bring adapter offline

DESCRIPTION: The FC driver has received a request to bring the adapter offline. This may occur when running lputil.

DATA: None

SEVERITY: Warning

LOG: LOG_INIT Verbose

ACTION: None required.

0462 INe: Too many cmd / rsp ring entries in SLI2 SLIM

DESCRIPTION: The configuration parameter for Scan-down is out of range.

DATA: (1) totiocb (2) MAX_SLI2_IOCB

SEVERITY: Error

LOG: Always

ACTION: This is a software driver error. If this problem persists, report these errors to Technical Support.

IP Traffic History (0600 - 0699)

0600 IPi: FARP-RSP received from DID <did>

DESCRIPTION: A FARP-ELS command response was received.

DATA: None

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: None required, informational.

0601 IPi: FARP-REQ received from DID <did>

DESCRIPTION: A FARP-ELS command response was received.

DATA: None

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: None required, informational.

0602 IPw: IP Response Ring <num> out of posted buffers

DESCRIPTION: The IP ring returned all posted buffers to the driver and is waiting for the driver to post new buffers. This could mean the host system is out of TCP/IP buffers.

DATA: (1) fc_missbufcnt (2) NoRcvBuf

SEVERITY: Warning

LOG: LOG_IP Verbose

ACTION: Try allocating more IP buffers (STREAMS buffers or mbufs) of size 4096 and/or increasing the post-ip-buf lpfc configuration parameter, then reboot the system.

0603 IPw: xmi Sequence completion error

DESCRIPTION: A XMIT_SEQUENCE command completed with a status error in the IOCB.

DATA: (1) ulpStatus (2) ulploTag (3) ulpWord[4] (4) did

SEVERITY: Warning

LOG: LOG_IP Verbose

ACTION: If there are many errors to one device, check the state of the remote PortID. The driver attempts to recover by creating a new exchange to the remote driver.

0605 IPw: No room on IP xmit queue

DESCRIPTION: The system is generating the IOCB commands to be processed faster than the HBA can process them.

DATA: (1) xmitnroom

SEVERITY: Warning

LOG: LOG_IP Verbose

ACTION: Check the state of the link. If the link is up and running, reconfigure the xmit queue size to be larger. Note, a larger queue size may require more physical connections to the Fibre Channel network.

0606 IPi: XRI Create for IP traffic to DID <DID>

DESCRIPTION: The lpfc driver is missing an exchange resource identifier (XRI) for this node and needs to create one prior to the transmit operation.

DATA: None

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: None required, informational.

0607 IPi: XRI response from DID with XRI <xri> and status <ulpStatus>

DESCRIPTION: The driver received an XRI response from SLI with the resulting exchange resource ID and status.

DATA: None

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: None required, informational.

0608 IPw: IP packet timed out

DESCRIPTION: An IP IOCB command was posted to a ring and did not complete within the timeout seconds.

DATA: None

SEVERITY: Warning

LOG: LOG_IP Verbose

ACTION: If no IP packet is going through the HBA, reboot the system. If the problem persists, contact Technical Support.

0609 IPe: Network buffer and DMA address mismatch

DESCRIPTION: An IP buffer free operation found a mismatch between an IP buffer and its dma address.

DATA: (1) pib (2) ip buff found (3) ip buf actual (4) dma address

SEVERITY: Error

LOG: Always

ACTION: Stop traffic and reboot the system.

0610 IPi: FARP Request sent to remote DID

DESCRIPTION: A send to a remote IP address has no node in the driver's nodelists. Send a FARP request to obtain the node's HW address.

DATA: (1) IEEE[0] (2) IEEE[1] (3) IEEE [2] (4) IEEE [3] (5) IEEE [4] (6) IEEE [5]

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: Issue FARP and wait for PLOGI from remote node.

0611 IPe: Dropping received IP packet

DESCRIPTION: The driver decided to drop an IP Response ring entry.

DATA: (1) ulpStatus (2) ulpWord

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver or firmware problem. If the problem persists, report the error to Technical Support.

FCP Traffic History (0700 - 0799)

0701 FPi: Issue Abort Task Set to TGT <num> LUN <num>

DESCRIPTION: The SCSI layer detected that it needs to abort all I/O to a specific device. This causes the FCP Task Management command to abort the I/O in progress.

DATA: (1) rpi (2) flags

SEVERITY: Information

LOG: LOG_FCP Verbose

ACTION: Check the state of the device in question.

0702 FPi: Issue Target Reset to TGT <num>

DESCRIPTION: The SCSI layer detected that it needs to abort all I/O to a specific target. This results in an FCP Task Management command to abort the I/O in progress.

DATA: (1) rpi (2) flags

SEVERITY: Information

LOG: LOG_FCP Verbose

ACTION: Check the state of the target in question.

0703 FPi: Issue LUN Reset to TGT <num> LUN <num>

DESCRIPTION: The SCSI layer detected that it must abort all I/O to a specific device. This results in an FCP Task Management command to abort the I/O in progress.

DATA: (1) rpi (2) flags

SEVERITY: Information

LOG: LOG_FCP Verbose

ACTION: Check the state of the device in question.

0710 FPi: Iodone <target>/<lun> error <result> SNS <lp> <lp3>

DESCRIPTION: This error indicates that the Fibre Channel driver is returning a SCSI command to the SCSI layer in error or with sense data.

DATA: (1) retry (2) resid

SEVERITY: Information

LOG: LOG_FCP Verbose

ACTION: None required, informational.

0712 FPe: SCSI layer issued abort device

DESCRIPTION: The SCSI layer is requesting the driver to abort I/O to a specific device.

DATA: (1) target (2) lun (3)

SEVERITY: Error

LOG: Always

ACTION: Check the state of the device in question.

0713 FPe: SCSI layer issued target reset

DESCRIPTION: The SCSI layer is requesting the driver to abort I/O to a specific target.

DATA: (1) target (2) lun

SEVERITY: Error

LOG: Always

ACTION: Check the state of the target in question.

0714 FPe: SCSI layer issued bus reset

DESCRIPTION: The SCSI layer is requesting the driver to abort all I/Os to all targets on this HBA.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: Check the state of the targets in question.

0716 FPi: FCP Read Underrun, expected <len>, residual <resid>

DESCRIPTION: An FCP device provided less data than was requested.

DATA: (1) fcpi_parm (2) cmnd[0] (3) underflow

SEVERITY: Information

LOG: LOG_FCP Verbose

ACTION: None required, informational.

0721 FPi: Cmpl Inquiry SN on NPort <nlp_DID>

DESCRIPTION: An INQUIRY Serial Number (page x83) completed. This information is saved by the driver.

DATA: (1) ulpStatus (2) rspStatus2 (3) rspStatus3 (4) scsi_id (5) lun_id

SEVERITY: Information

LOG: Log_FCP Verbose

ACTION: None required, informational.

0722 FPi: INQUIRY SN info

DESCRIPTION: This is the serial number of the device that will be saved.

DATA: (1) datap* (2) datap + 3* (3) datap + 7 (4) rspResId

SEVERITY: Information

LOG: Log_FCP Verbose

ACTION: None required, informational.

0723 FPi: Issue Inquiry SN on NPort <nlp_DID>

DESCRIPTION: Issuing an INQUIRY Serial Number (page x83) FCP command.

DATA: (1) failMask (2) scsi_id (3) lun_id

SEVERITY: Information

LOG: Log_FCP Verbose

ACTION: None required, informational.

0724 FPi: Issue Inquiry PO on NPort <nlp_DID>

DESCRIPTION: N/A

DATA: Issuing an INQUIRY Serial Number (page x0) FCP command.

SEVERITY: Information

LOG: Log_FCP Verbose

ACTION: None required, informational.

0725 FPe: Inquiry Serial Number: invalid length

DESCRIPTION: An INQUIRY SN command completed with an invalid serial number length

DATA: (1) sizasn (2) scsi_id (3) scis_id (4) lun_id

SEVERITY: Error

LOG: Always

ACTION: Check remote NPORT for potential problem.

0726 FPe: INQUIRY SN cmd failed

DESCRIPTION: The INQUIRY Serial Number (page x83) failed.

DATA: (1) ulpStatus (2) fcpi_parm (3) m_target (4) m_lun

SEVERITY: Error

LOG: Always

ACTION: Check if target device supports this command.

0727 FPi: Cmpl INQUIRY PO on NPort <nlp_DID>

DESCRIPTION: An INQUIRY (page 0) completed. This information is saved by the driver.

DATA: (1) ulpStatus (2) rspStatus2 (3) rspStatus3 (4) scsi_id (5) lun_id

SEVERITY: Information

LOG: Log_FCP Verbose

ACTION: None required, informational.

0728 FPe: INQUIRY Page 0 cmd failed

DESCRIPTION: The INQUIRY (page 0) failed.

DATA: (1) ulpStatus (2) fcpi_parm (3) scsi_id (4) lun_id

SEVERITY: Error

LOG: Always

ACTION: Check if target device supports this command.

0729 FPw: FCP cmd <cmnd> failed <target>/<lun>

DESCRIPTION: The specified device failed an FCP command.

DATA: (1) status (2) result (3) xri (4) iotag

SEVERITY: Warning

LOG: LOG_FCP Verbose

ACTION: Check the state of the target in question.

0730 FPw: FCP command failed: RSP

DESCRIPTION: The FCP command failed with a response error.

DATA: (1) Status2 (2) Status3 (3) ResId (4) SnsLen (5) RspLen (6) Info3

SEVERITY: Warning

LOG: LOG_FCP Verbose

ACTION: Check the state of the target in question.

0734 FPw: FCP Read Check Error

DESCRIPTION: The issued FCP command returned a read check error.

DATA: (1) fcpDI (2) rspResId (3) fcpi_parm (4) cdb[0]

SEVERITY: Warning

LOG: LOG_FCP Verbose

ACTION: Check the state of the target in question.

0735 FPw: FCP Read Check Error with Check Condition

DESCRIPTION: The issued FCP command returned a read check error and a check condition.

DATA: (1) fcpDI (2) rspResId (3) fcpi_parm (4) cdb[0]

SEVERITY: Warning

LOG: LOG_FCP Verbose

ACTION: Check the state of the target in question.

0736 FPi: Received Queue Full status from FCP device <tgt> <lun>

DESCRIPTION: Received a Queue Full error status from specified FCP device.

DATA: (1) qfull_retry_count (2) qfull_retries (3) currentOutstanding (4) maxOutstanding

SEVERITY: Information

LOG: LOG_FCP Verbose

ACTION: None required, informational.

0737: <ASC ASCQ> Check condition received

DESCRIPTION: The issued FCP command resulted in a check condition.
DATA: (1) CFG_DELAY_RSP_ERR (3) *lp
SEVERITY: Information
LOG: LOG_FCP | LOG_CHK_COND Verbose
ACTION: None required, informational.

0747 FPi: Cmpl target reset

DESCRIPTION: A driver-initiated target reset completed.
DATA: (1) scsi_id (2) lun_id (3) Error (4) statLocalError (5) *cmd + WD7
SEVERITY: Information
LOG: LOG_FCP Verbose
ACTION: None required, informational.

0748 FPi: Cmpl LUN reset

DESCRIPTION: A driver-initiated LUN reset completed.
DATA: (1) scsi_id (2) lun_id (3) Error (4) statLocalError (5) *cmd + WD7
SEVERITY: Information
LOG: LOG_FCP Verbose
ACTION: None required, informational.

0749 FPi: Cmpl Abort Task Set

DESCRIPTION: A driver-initiated abort task set completed.
DATA: (1) scsi_id (2) lun_id (3) Error (4) statLocalError (5) *cmd + WD7
SEVERITY: Information
LOG: LOG_FCP Verbose
ACTION: None required.

0753 FPe: Inquiry Serial Number: invalid length

DESCRIPTION: An INQUIRY SN command completed with an invalid serial number length.
DATA: (1) sizeSN (2) j (3) scsi_id (4) lun_id
SEVERITY: Error
LOG: Always
ACTION: Check state of target in question.

0754 FPe: SCSI timeout

DESCRIPTION: An FCP IOCB command was posted to a ring and did not complete within ULP timeout seconds.
DATA:(1) did (2) sid (3) lun (4) command (5) iotag
SEVERITY: Error
LOG: Always
ACTION: If I/O is not going through the adapter, reboot the system; otherwise check the state of the target in question. If the problem persists, contact Technical Support.

Node Table Events (0900 - 0999)

0900 NDi: Cleanup node for NPort <nlp_DID>

DESCRIPTION: The driver node table entry for a remote NPort was removed.

DATA: (1) nlp_flag (2) nlp_state (3) nlp_rpi

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0901 NDi: FIND node DID mapped

DESCRIPTION: The driver is searching for a node table entry, on the mapped node list, based on the DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0902 NDi: FIND node DID mapped

DESCRIPTION: The driver is searching for a node table entry, on the mapped node list, based on DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0903 NDi: Add scsiid <sid> to BIND list

DESCRIPTION: The driver is putting the node table entry on the binding list.

DATA: (1) bind_cnt (2) nlp_DID (3) bind_type (4) blp

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0904 NDi: Add NPort <did> to <list> list

DESCRIPTION: The driver is moving the node table entry on the specified list.

DATA: (1) nlp_flag (2) blp

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0910 NDi: FIND node DID unmapped

DESCRIPTION: The driver is searching for a node table entry on the unmapped node list, based on DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0911 NDi: FIND node DID unmapped

DESCRIPTION: The driver is searching for a node table entry, on the unmapped node list, based on DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0917 NDi: PUT END nodelist

DESCRIPTION: The driver is freeing a node table entry buffer.

DATA: (1) bp (2) fc_free

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0927 NDi: GET nodelist

DESCRIPTION: The driver is allocating a buffer to hold a node table entry.

DATA: (1) bp (2) fc_free

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0928 NDi: PUT nodelist

DESCRIPTION: The driver is freeing a node table entry buffer.

DATA: (1) bp (2) fc_free

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0929 NDi: FIND UNMAPPED node DID

DESCRIPTION: The driver is searching for a node table entry, on the unmapped node list, based on DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0930 NDi: FIND MAPPED node DID

DESCRIPTION: The driver is searching for a node table entry, on the mapped node list, based on DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0931 NDi: FIND PLOGI node DID

DESCRIPTION: The driver is searching for a node table entry, on the PLOGI node list, based on DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0932 NDi: FIND REGLOGIN node DID

DESCRIPTION: The driver was searching for a node table entry on the REGLOGIN node list, based on the DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0933 NDi: FIND PRLI node DID

DESCRIPTION: The driver was searching for a node table entry on the PRLI node list, based on the DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0934 NDi: FIND ADISC node DID

DESCRIPTION: The driver was searching for a node table entry on the ADISC list, based on the DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0935 NDi: FIND NPR node DID

DESCRIPTION: The driver was searching for a node table entry on the NPR node list, based on the DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0936 NDi: FIND UNUSED node DID

DESCRIPTION: The driver was searching for a node table entry on the UNUSED node list, based on the DID.

DATA: (1) nlp (2) nlp_DID (3) nlp_flag (4) data1

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

0937 NDi: FIND node DID <did> NOT FOUND

DESCRIPTION: The driver was searching for a node table entry based on the DID and the entry was not found.

DATA: (1) order

SEVERITY: Information

LOG: LOG_NODE Verbose

ACTION: None required, informational.

Miscellaneous Events (1200 - 1299)

1200: Cannot unload driver while lpfcdiag interface is active

DESCRIPTION: An attempt was made to unload the driver while the DFC interface was active.

DATA: (1) refcnt (2) instance

SEVERITY: Error

LOG: Always

ACTION: Exit any application that uses the DFC diagnostic interface before attempting to unload the driver.

1208 Mli: C_CT request error

DESCRIPTION: The CT response returned more data than the user buffer could hold.

DATA: (1) dfc_flag (2) 4096

SEVERITY: Information

LOG: LOG_MISC Verbose

ACTION: Modify the user application issuing a CT request to allow for a larger response buffer.

1210 Mle: Convert ASC to hex. Input byte cnt <1

DESCRIPTION: ASCII string to hexadecimal conversion failed. The input byte count is greater than 1.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver problem. If the problem persists, report the error to Technical Support.

1211 Mle: Convert ASC to hex. Input byte cnt > max <num>

DESCRIPTION: ASCII string to hexadecimal conversion failed. The input byte count exceeds max <num>.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver problem. If problems persist report, these errors to Technical Support.

1212 Mle: Convert ASC to hex. Output buffer too small

DESCRIPTION: ASCII string to hexadecimal conversion failed. The output buffer byte size is less than 1/2 of the input byte count. Every two input characters (bytes) require one output byte.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver problem. If the problem persists, report the error to Technical Support.

1213 Mlc: Convert ASC to hex. Input char seq not ASC hex

DESCRIPTION: The ASCII hexadecimal input string contains a non-ASCII hex character.

DATA: None

SEVERITY: Error configuration

LOG: Always

ACTION: Make necessary changes to the lpfc configuration file.

1214 Mle: Cannot unload driver, IP interface still attached.

DESCRIPTION: An attempt was made to unload the driver while the IP interface was active.

DATA: (1) refcnt (2) brd_no

SEVERITY: Error

LOG: Always

ACTION: Exit the IP interface before attempting to unload the driver.

Link Events (1300 - 1399)

1300 LKe: Re-establishing Link, timer expired

DESCRIPTION: The driver detected a condition where it had to re-initialize the link.

DATA: (1) fc_flag (2) fc_ffstate

SEVERITY: Error

LOG: Always

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

1301 LKi: Re-establishing Link

DESCRIPTION: The driver detected a condition in which it had to re-initialize the link.

DATA: (1) status (2) status1 (3) status2

SEVERITY: Information

LOG: LOG_LINK_EVENT Verbose

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

1302 Lli: Reset link speed to auto. 1G HBA cfg'd for 2G

DESCRIPTION: The driver is re initializing the link speed to auto-detect.

DATA: (1) current link speed

SEVERITY: Warning

LOG: LOG_LINK_EVENT Verbose

ACTION: None required, informational.

1303 LKe: Link Up Event <eventTag> received

DESCRIPTION: A link up event was received. It is also possible for multiple link events to be received together.

DATA:(1) fc_eventTag (2) granted_AL_PA (3) UlnkSpeed (4) alpa_map[0]

Detail: If link events received, log (1) last event number received, (2) ALPA granted, (3) Link speed (4) number of entries in the loop init LILP ALPA map. An ALPA map message is also recorded if LINK_EVENT Verbose mode is set. Each ALPA map message contains 16 ALPAs.

SEVERITY: Error

LOG: Always

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

1304 LKw: Link Up Event ALPA map

DESCRIPTION: A link up event was received.

DATA: (1) wd1 (2) wd2 (3) wd3 (4) wd4

SEVERITY: Warning

LOG: LOG_LINK_EVENT Verbose

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

1305 LKe: Link Down Event <eventTag> received

DESCRIPTION: A link down event was received.

DATA: (1) fc_eventTag (2) hba_state (3) fc_flag

SEVERITY: Error

LOG: Always

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network.

1306 LKe: Link Down timeout on HBA

DESCRIPTION: The link was down for more than the configured link-down-tmo seconds.

DATA: (1) hba_state (2) fc_flag (3) fc_ns_retry

SEVERITY: Warning

LOG: LOG_LINK_EVENT | LOG_DISCOVERY verbose

ACTION: Check HBA cable/connections to Fibre Channel network.

1307 LKi: READ_LA mbox error <mbxStatus> state <hba_state>

DESCRIPTION: The driver cannot determine what type of link event occurred.

DATA: None

SEVERITY: Information

LOG: LOG_LINK_EVENT Verbose

ACTION: If numerous link events are occurring, check the physical connections to the Fibre Channel network. May indicate a possible hardware or firmware problem.

1308 LKe: Out of memory; Ignoring the link attention.

DESCRIPTION: The driver is unable to process the link attention due to insufficient memory.

DATA: None

SEVERITY: Error

LOG: Always

ACTION: This error could indicate a software driver problem. If the problem persists, contact Technical Support.

Log Messages - IOCTL Events (1600 - 1699)

1600 IOi: dfc_ioctl entry

DESCRIPTION: The entry point for processing diagnostic ioctl.

DATA:(1) c_cmd (2) c_arg1 (3) c_arg2 (4) c_outsz

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: None required, informational.

1601 IOi: dfc_ioctl exit

DESCRIPTION: The exit point for processing diagnostic ioctl.

DATA: (1) rc (2) c_outsz (3) c_dataout

SEVERITY: Information

LOG: LOG_IP Verbose

ACTION: None required, informational.

1602 IOi: dfc_data_alloc

DESCRIPTION: Allocating data buffer to process dfc ioctl.

DATA: (1) fc_dataout (2) fc_outsz

SEVERITY: Information

LOG: LOG_IOC Verbose

ACTION: None required, informational.

1603 IOi: dfc_data_free

DESCRIPTION: The data buffer is being freed to process dfc ioctl.

DATA: (1) fc_dataout (2) fc_outsz

SEVERITY: Information

LOG: LOG_IOC Verbose

ACTION: None required, informational.

1604 IOe: lpfc_ioctl:error

DESCRIPTION: The SCSI send request buffer size limit was exceeded.

DATA: (1) error number index

SEVERITY: Error

LOG: Always

ACTION: Reduce the application program's SCSI send request buffer size to less than 320K bytes.