



Configuring the Driver for Solaris

Version 5.02f (x86)

Table of Contents

[Introduction](#)

[List Adapters](#)

[View Adapter Information](#)

[Understand Device Numbering](#)

[Set Up Persistent Binding](#)

[Add New Targets](#)

[Change Parameters or Bindings](#)

[Set Up Target/LUN Blocking](#)

[Enable IP Networking](#)

[Enable IP Networking for Multiple Adapters](#)

[Create Additional Hostname.lpf# Files](#)

[Probe for FCP Disk Targets](#)

[Probe for FCP Tape Targets](#)

[Configure the System-Wide File](#)

[View lpfc.conf Table](#)

Introduction

You can configure the driver for Solaris by using two primary methods:

- Using the LightPulse Utility (lputil)
- Using a text editor to edit configuration files

Using lputil

lputil allows you to view adapter information and set up (maintain) persistent binding. Start the utility by entering the following command:

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

The Main menu displays. The View Adapter Information and Set Up Persistent Binding topics are included in the Configuration Table of Contents. For additional information about lputil, see the online Utility Reference manual.

Editing Configuration Files

Use your text editor to change the parameters for the following files.

- **lpfc.conf file** - Specifies all tunable parameters for all host adapters attached to the host system. This file is located in the driver installation directory (typically / kernel/drv) and contains all the variables that control driver initialization.
- **sd.conf file** - Configure only when the driver is configured for storage. sd.conf tells the system which disks to probe when the system is booted. For FCP targets with multiple LUNs or target IDs greater than 17, sd.conf must be modified. Also, edit this file to limit the SCSI layer probing of targets and LUNs on an adapter by adapter basis.
- **hostname.lpfcX file** - Change parameters in this file to enable networking on two or more installed host adapters in the same system.
- **system-wide files** - Set moddir if using a non-standard installation directory, set forcload if enabling IP networking.

List Adapters

When you start the LightPulse Utility (lputil), all adapters are listed with information such as the host adapter number, instance number (i.e. lpfc0), board model type, whether it is a standard or non-standard device and whether the adapter is ready to use.

From the Main menu, enter 1, List Adapters.

MAIN MENU

- 1. List Adapters
- 2. Adapter Information
- 3. Firmware Maintenance
- 4. Reset Adapter
- 5. Persistent Binding

- 0. Exit

Enter choice => 1

View Adapter Information

The LightPulse Utility (lputil) allows you to view information for a selected adapter. Once you start lputil, the Main menu opens:

MAIN MENU

1. List Adapters
 2. Adapter Information
 3. Firmware Maintenance
 4. Reset Adapter
 5. Persistent Binding
 0. Exit
- Enter choice => 2

Select choice #2. If you have multiple adapters, a list displays, you select an adapter and the Adapter Information Menu opens. If you have only one adapter, the Adapter Information Menu opens for that adapter.

ADAPTER INFORMATION MENU

1. PCI Configuration Parameters
 2. Adapter Revision Levels
 3. Wakeup Parameters
 4. IEEE Address
 5. Loop Map
 6. Status and Counters
 7. Link Status
 8. Configuration Parameters
 0. Return to Main Menu
- Enter choice =>

PCI Configuration Parameters - Displays 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 - Displays firmware revision levels, including kernel and overlay version information.

Wakeup Parameters - Displays BIOS status and version, as well as SLI (service level interface).

IEEE Address - Displays 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 - Tracks such activities as link failure, loss synch and elastic overlay.

Configuration Parameters - Displays parameters such as link failure count, loss sync count and timeout.

Understand 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 gets 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 get assigned, they will stay the same between reboots if there are no LightPulse hardware changes to the system.

Set Up Persistent Binding

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 start the lputil utility, type:

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

The Main menu opens:

MAIN MENU

1. List Adapters
2. Adapter Information
3. Firmware Maintenance
4. Reset Adapter
5. Persistent Binding
0. Exit

Enter choice => 5

Select choice #5. The Persistent Binding menu opens:

PERSISTENT BINDINGS MENU

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

0. Return to Main Menu

Enter choice => 4

Use this menu to view bindings and nodes, establish new bindings, or delete bindings.

1. To establish new bindings, select option => 4.
2. Select an adapter.
3. Select 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

4. Enter the Node Name, Port Name, or D_ID, depending on binding method.
5. Enter the target number => 0.

Add New Targets

You can perform on-the-fly configuration changes in Solaris 9, without rebooting, in the following way:

- 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;  
.  
.  
.
```

Save the file and exit vi.

- Quiesce all I/O on the device.
- Force a reread of the file with the Solaris command **update_drv -f sd**.

Change 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.

Set Up Target/LUN Blocking

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).

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

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 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.

Enable IP Networking

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.

1. If IP networking support was not enabled during initial driver installation, update the `lpfc.conf` file to enable IP networking.
2. Any name given in the `hostname.lpfc<X>` files must have a valid IP address associated with it. 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 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.

Update the `lpfc.conf` file

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

```
vi /kernel/drv/lpfc.conf
```

2. Enter a binary 1 for the network-on variable to enable IP networking (1 = on, 0 = off).

Modify the /etc/hosts file

The hosts file associates an IP address to a host adapter's network host name.

1. Use a text editor to edit the /etc/hosts file.

```
vi /etc/hosts
```

Use the format **IP_address host_name**

where **IP_address** is the host's IP address and **host_name** is the name of the host.

Example:

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

Create a Hostname File for Each Adapter

The network name identifies the host adapter on a Fibre Channel IP network and is associated with a unique IP address. The network name's IP address is configured in the /etc/hosts file.

1. Create a hostname.lpfcX file for each IP networking host adapter, where X is the host adapter's instance number in decimal format.

```
vi /etc/hostname.lpfcX
```

2. Reboot the system.

Verify 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.

Enable IP Networking for Multiple Adapters

To enable networking on two or more installed host adapters in the same system, additional `hostname.lpf<X>` files need to be created manually (where **X** is the instance number of the host adapter being configured). Upon rebooting, Solaris will automatically configure and setup 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.lpf<X>` files must have a valid IP address associated with it. The IP address is established through NIS, DNS or other addressing scheme.

Alternately, re-run the installation procedure (doing so is called an update install). This provides the opportunity to enable IP networking support, which then updates all the configuration file settings. Reboot the system for these changes to take effect.

Create 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`.

Probing for FCP Disk Targets -- Configuring sd.conf

The following discussion applies only when you have enabled FCP functionality in lpfc. 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.conf for Solaris 7, Solaris 8 and Solaris 9. The configuration file 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

The following discussion applies only when you have enabled FCP functionality in lpfc. 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 is usually found in /kernel/drv. See the driver.conf(4) man page for additional details. The tape target driver is called st.conf, 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.

Note 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.

Configure 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 variables are in lpfc.conf and the target driver configuration file. Listed below are a few cases which may require modification of the /etc/system.

Driver installed in non-standard directory (moddir)

The moddir variable 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 variable 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 variable 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 variable forceload must be set to drv/clone for correct operation. The lpfn driver is dependent upon 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.

View the lpfc.conf Configuration Table

This file is located in the driver installation directory (typically /kernel/drv) and contains all the variables that control driver initialization. The table is color coded based on the effect of the parameter.

Note To set these values per adapter, append **lpfcX-** to the configuration parameter (where **X** is the driver instance number) for example, setting **lpfc0-num-iocbs** to 1000 allocates 1000 I/O control blocks for the driver interface **lpfc0**.

A clear row indicates that the parameter affects both FCP storage and IP networking

A green row indicates that the parameter affects FCP storage only.

A yellow row indicates that the parameter affects IP networking only.

Controller-Specific and Global Parameters

Parameter values are set per host bus adapter. For example, setting **numb-iocbs** to 1000 allocates 1000 I/O control blocks for every host bus adapter present in the host. If you want to have configuration parameters apply to a specific host bus adapter, prepend **lpfcX-** to the configuration parameter (where **X** is the driver instance number). For example, setting **lpfc0-num-iocbs** to 2000 allocates 2000 I/O control blocks for the driver interfaces **lpfc0**.

All variables indicated as controller-specific (C) in the following table are global-specific unless prepended with the **lpfcX-** code. Those variables indicated as global (G) in the following table are never controller-specific, they are only global.

Click on a parameter to view a detailed description.

Variable	Controller - or Global - Specific	Default	Min	Max	Comments
ack0	C	0	0	1	Use ACK0 for Class 2.
automap	C	2	0	3	Automatically assigns SCSI IDs to FCP targets detected.
check-cond-err	C	0	Boolean 0=false, 1=true		Treat certain FCP check conditions as FCP RSP errors.
cr-delay	C	0	0	63	Specify a count of milliseconds after which an interrupt response is generated if cr-count has not been satisfied.
cr-count	C	8	2	255	Specify a count of I/O completions after which an interrupt response is generated.
delay-rsp-err	C	0	0	1	Treat FCP RSP errs like no-device-delay.
dqfull-throttle	C	1	Boolean 0=false, 1=true		Decrement LUN throttle on queue fulls.
dqfull-throttle-up-inc	C	1	0	128	Amount to increment LUN queue depth each time.
dqfull-throttle-up-time	C	1	0	30	Time interval, in secs, to increment LUN queue depth.
fcp-bind-WWNN	G	inactive	N/A	N/A	Bind specific SCSI IDs to targets based on Fibre Channel WWNN.

fcplib-bind-WWPN	G	inactive	N/A	N/A	Bind specific SCSI IDs to targets based on Fibre Channel WWPN.
fcplib-bind-DID	G	inactive	N/A	N/A	Bind specific SCSI IDs to targets based on Fibre Channel port ID.
fcplib-class	C	3	2	3	Fibre Channel class for FCP data transmission.
fcplibfabric-tmo	C	0	0	255	Extra FCP cmd timeout when connected to a fabric (in seconds).
fcplib-on	C	1	Boolean 0=false, 1=true		Turn on FCP.
fdmi-on	C	0	Boolean 0=false, 1=true		Turn on FDMI.
ip-class	C	3	2	3	Fibre Channel class for IP data transmission.
link-speed	C	0	0	2	Set link speed: auto, 1G, 2G
linkdown-tmo	C	30	0	255	(seconds) How long driver waits before deciding FC link is down.
lpfcNtM-lun-throttle	C	none	1	128	Override lun-queue-depth for lpfc adapter N target M.
lpfcNtM-tgt-	C	none	1	10240	Override tgt-queue-depth for lpfc adapter N target M.
lun-queue-depth	G	30	1	128	Default max commands sent to a single logical unit (disk).
network-on	C	0	Boolean 0=false, 1=true		Turn on IP networking.

no-device-delay	G	1	0	30	Delay to fail back an I/O in seconds.
nodev-holdio	C	0	0	1	Hold I/O errors forever if device disappears.
nodev-tmo	C	0	0	255	Seconds to hold I/O if device disappears.
num-bufs	C	1024	64	4,096	ELS and misc data buffers.
num-iocbs	C	1024	128	10,240	Adapter command/response structures.
scan-down	C	2	0	2	Select method for ALPA to SCSI ID mapping.
tgt-queue-depth	G	0	0	10240	Default max commands sent to a single target.
topology	C	0	0	6	Fibre Channel link topology (defaults to attempt loop first then point-to-point).
use-adisc	C	0	0	1	(Boolean)Send ADISC instead of PLOGI for device discovery or RSCN.
log-only	C	1	0	1	Logging to log file only.
log-verbose	C	0	0	0xff	(bit mask) Extra activity logging.
post-ip-buf	C	128	64	1024	Buffers to post to IP ring.
xmt-que-size	C	256	128	10,240	Transmit queue length for mbufs.

Detail of lpfc.conf Parameters

ack0

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.

automap

If automap is non-zero, SCSI IDs for all FCP nodes without persistent bindings will be automatically generated. 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 one of the FCP binding methods below is specified, then automap devices will use the same mapping method to try to preserve SCSI IDs between link down and link up. If no persistent binding method is specified, the binding is by WWNN for automap 1 (default), by WWPN for automap 2, and by D_ID for automap 3. If automap is 0, only devices with persistent bindings will be recognized by the system.

check-cond-err

The driver will treat check condition 0x44, Internal Target Failure, and check condition 0x20, Invalid Command Operation Code, as FCP RSP errors.

cr-delay

This value specifies a count of milliseconds after which an interrupt response is generated if cr-count has not been satisfied. This value is set to 0 to disable the Coalesce Response feature as default.

cr-count

This value specifies a count of I/O completions after which an interrupt response is generated. Default is set to 8.

delay-rsp-err

The driver will delay FCP RSP errors being returned to the upper SCSI layer based on the no-device-delay configuration parameter.

dqfull-throttle

dqfull-throttle-up-inc

dqfull-throttle-up-time

The variable dqfull-throttle causes lpfc to decrement a LUNs throttle (maximum 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 restore the queue depth back to the original.

The dqfull-throttle-up-time variable defines a time, in seconds, that will indicate 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, this variable increments the current queue depth by invoking the dqfull-throttle-up-inc variable (does not exceed the lun-queue-depth). If both parameters are set to 1, the driver increments the current queue depth once per second until it reaches the lun-queue-depth. The only other way to restore the queue depth back to the original LUN throttle (besides rebooting), 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.

fcp-bind-WWNN

Set up persistent FCP bindings based on a target device's WWNN. It is a global parameter across all LightPulse host bus adapters in the system. This binding guarantees that target assignments will be preserved between reboots. The format for a bind entry is "**NNNNNNNNNNNNNNNNNN**:lpfc*XtY*" where **NNNNNNNNNNNNNNNNNN** is a 16-digit representation of the targets World-Wide Node Name, *X* is the physical hardware adapter 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.

Example:

```
fcp-bind-WWNN="20000020370c396f:lpfc1t0","20000020370c27f7:lpfc0t2";
```

fcp-bind-WWPN

Set up persistent FCP bindings based on a target device's WWPN. It is a global parameter across all LightPulse host bus adapters in the system. This binding guarantees that target assignments will be preserved between reboots. The format for a bind entry is "**NNNNNNNNNNNNNNNNNN**:lpfc**XtY**", where **NNNNNNNNNNNNNNNNNN** is a 16-digit representation of the target's World-Wide Port Name, **X** is the physical hardware adapter 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.

Example:

```
fcp-bind-WWPN="21000020370cf8263:lpfc1t0";
```

fcp-bind-DID

Set up persistent FCP bindings based on a target device's Port ID. It is a global parameter across all LightPulse host bus adapters in the system. This binding guarantees that target assignments will be preserved between reboots. The format for a bind entry is "**NNNNNN**:lpfc**XtY**", where **NNNNNN** is a 6-digit representation of the target's Port ID, **X** is the physical hardware adapter 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.

Example:

```
fcp-bind-DID="0000ef:lpfc0t0";
```

fcp-class

The driver is capable of transmitting FCP data in Class 2 or Class 3. This parameter defaults to Class 3 transmission.

fcp-fabric-tmo

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.

fcp-on

Enable the FCP functionality of the driver. This variable is Boolean; zero disables FCP, while non-zero enables it.

fdmi-on

Enable the Fabric Device Management Interface (FDMI) functionality of Ipfc. FDMI enables the management of devices such as HBAs through the fabric. This variable is Boolean; zero disables FDMI, while non-zero enables it.

ip-class

Fibre Channel is capable of transmitting IP data in Class 2 or Class 3. The driver defaults to using Class 3 transmission. Change this variable only if you have a good understanding of the properties of Fibre Channel and a strong need to use one of the other two classes.

link-speed

This causes the driver to set the adapter link speed to auto-detect (0), 1 gigabits/sec (1) or 2 gigabits/sec (2). This setting applies only to the LightPulse 2 gigabits/sec HBA family of adapters.

linkdown-tmo

How long the driver will wait (1 to 255 seconds) before declaring the Fibre Channel link to be down. Commands waiting to be serviced will be aborted. For instance, FCP commands will be returned to the target driver. The lower this value is, the faster the driver will fail commands back to the upper levels. There is a tradeoff here: small values risk retrying commands when the link is bouncing, while large values risk delaying failover in a redundant loop environment. Setting linkdown-tmo to 0 will cause all commands issued to the device driver to be held internally (never failed) until the link comes back up.

lun-queue-depth

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. A single logical unit will never be sent more commands than allowed by lun-queue-depth; however, fewer may be sent when sd-max-throttle or tgt-queue-depth is reached for the entire target.

lpfcNtM-tgt-throttle

The maximum number of outstanding commands to permit for any logical unit for any target including all LUNs on that target. This value overrides tgt-queue-depth.

lpfcNtM-lun-throttle

The maximum number of outstanding commands to permit for any logical unit on a specific target. This value overrides lun-queue-depth.

network-on

This variable controls whether the driver 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.

no-device-delay

This variable [0 to 30] determines the length of the interval between deciding to fail an I/O because there is no way to communicate with its particular device (for example, 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

This variable controls whether I/O errors are held by the driver if a device on the loop disappears. If set, I/O errors will be held until the device returns to the loop (potentially indefinitely). The upper sd layer may retry the command once the error is returned.

nodev-tmo

This variable controls how long I/O errors are held by the driver if a device on the loop disappears. If set, I/O errors will be held for the specified number of seconds. The upper sd layer may retry the command once the error is returned.

num-bufs

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 several FCP commands to the adapter, then 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. **Important:** The driver must always be configured with at least several dozen ELS command buffers; we recommend at least 256.

num-iocbs

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.

scan-down

Three scanning algorithms are used to map an AL_PA to a SCSI ID in a private loop. Refer to the "informative annex" of the FC-AL profile, which describes the algorithm for this translation. 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. If the adapter is not attached to a fabric, these parameters do not apply (use persistent binding).

If scan-down is 1, SCSI IDs are scanned starting from AL_PA 0xEF down through AL_PA 1 (for example, if only two targets exist on the loop, the one with the highest ALPA is assigned SCSI ID 0 and the other one is assigned SCSI ID 1).

If scan-down is 0, SCSI IDs are scanned starting from AL_PA 1 and up (for example, if only two targets exist on the loop, the one with the lowest AL_PA gets assigned SCSI ID 0 and the other one is assigned SCSI ID 1).

If scan-down is 2 (default value), SCSI IDs are hard coded based on the AL_PA (0xEF=SCSI ID 0, 0xE8=SCSI ID 1,...). This parameter is used to maintain backwards compatibility with the SCSI ID assignments of the 2.05 Solaris driver.

tgt-queue-depth

The driver uses this value as the default limit for the number of simultaneous commands to issue to a single target. 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, fewer may be sent when `sd-max-throttle` is reached for the entire target.

topology

This variable controls the Fibre Channel topology expected by the driver at boot time. Fibre Channel offers point-to-point, fabric and arbitrated loop topologies. This variable encodes these topologies as follows:

0 - Attempt loop mode, if it fails attempt point-to-point mode. (default)

2 - Attempt point-to-point mode only.

4 - Attempt loop mode only.

6 - Attempt point-to-point mode, if it fails attempt loop mode.

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) are supported in firmware version 1.00a2 for the LP98XX family of products, and in firmware version 3.20 and higher for all other LightPulse products.

use-adisc

This variable controls the ELS command used for FCP address authentication during discovery. If set, ADISC is used, otherwise, PLOGI is used.

log-only

When set to 1, log messages are logged only to syslog. When set to 0, log messages are also display on the console.

log-verbose

(bit mask) When set to non-zero this variable causes the driver to generate copious additional messages concerning the state of the driver and the I/O operations it carries out. These messages may go to the system console and/or the system log file, `/var/adm/messages`. See Error Codes Table in the Online Troubleshooting Manual for detailed information on the bit mask.

This variable specifies the number 4K STREAMS buffers to allocate and post to the Fibre Channel IP ring. Increase this setting for better IP performance under heavy loading.

xmt-queue-size

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.