



Host Settings Affected by FC Solaris Host Utilities

Fibre Channel Solaris Host Utilities changes to host settings

The Fibre Channel (FC) Solaris Host Utilities sometimes recommend that settings on the host be modified. These recommendations are based on research and internal testing done at NetApp. Testers use different settings before arriving at the recommendations.

Note

The values listed here apply to the versions of the Host Utilities 5.1 and earlier using the FC protocol.

The recommended values may differ based on:

- ◆ Whether you are using a clustered storage system or a single, stand-alone storage system. With clustered systems, the values may vary depending on your cluster failover mode (cfmode).
- ◆ Whether you are running the current version of the Host Utilities or an earlier version of the Attach Kit, such as FC Solaris Attach Kit 2.0.
- ◆ Which environment you are using. The Solaris Host Utilities support Veritas Dynamic Multipathing (DMP) environments with either Emulex LightPlus Fibre Channel drivers (LPFC) or Sun native drivers and Solaris native environments using MPxIO.

Note

To make it immediately clear which environment is being used, this document sometimes specifies “DMP” for the Veritas DMP environment and “MPxIO” for the Solaris native environment. In some cases, the commands you use vary depending on which drivers you are using. In those cases, both the environment and driver types are specified.

When you are using a Veritas DMP environment, NetApp recommends that in most cases you use the default values except for the restore daemon settings.

This document provides the following summaries of the suggested settings:

- ◆ [“Tools to use when modifying host settings”](#) on page 2
- ◆ [“Recommended settings for Veritas DMP with LPFC drivers”](#) on page 3
- ◆ [“Recommended settings for Veritas DMP with Sun native drivers”](#) on page 12
- ◆ [“Recommended Host Utilities settings for MPxIO environment”](#) on page 13
- ◆ [“Recommended settings when using Veritas VxVM and DMP”](#) on page 14

Tools to use when modifying host settings

Tools to modify host settings

The Solaris Host Utilities provide two tools that modify settings on the host when you are using the FC protocol.

Command	Description
basic_config	Depending on the system setup and options used, modifies parameters in the following files: ◆ Veritas with LPFC drivers ❖ /kernel/drv/lpfc.conf ❖ /etc/system ◆ Veritas DMP with Sun Native Drivers ❖ /kernel/drv/ssd.conf ◆ Native OS ❖ /kernel/drv/ssd.conf ❖ /kernel/drv/sd.conf Note— You can manually set these parameters instead of using basic_config, if you choose.
create_binding	(Veritas DMP with LPFC drivers) Modifies parameters in /kernel/drv/lpfc.conf to create persistent bindings. Note— You can use other tools, such as <code>lputil</code> or <code>HBAnyware</code>, to create persistent bindings.

You can use the `basic_config` command with all of the supported Solaris Host Utilities environments that use FC. The options that you supply with the command differ based on the environment and the drivers used with it. For details on using these commands, see the *Solaris Host Utilities Installation and Administration Guide*.

The `create_binding` program applies to the Veritas stack using LPFC drivers.

Recommended settings for Veritas DMP with LPFC drivers

basic_config sets values in lpfc.conf based on cfmode

(Veritas DMP with LPFC) The `basic_config` command changes the `linkdown-tmo` and `nodev-tmo` parameters in the `/kernel/drv/lpfc.conf` file based on the `cfmode` that your system is using. Storage controller `cfmodes` have different failover characteristics, namely the length of time it takes for the storage controller to completely fail-over. The `basic_config` command adjusts these values based on the length of time normally required for a host to recover from a failover based on its `cfmode`.

cfmode-dependent values for parameters on systems using DMP with LPFC drivers			
Parameter	Cfmode	Value	Description
linkdown-tmo	Single-image, Partner, Standby	30 seconds	The <code>basic_config</code> command resets the value to the driver default, which is 30 seconds. If the driver detects that the link to the HBA has disappeared, it waits 30 seconds for the link to reappear.
linkdown-tmo	Stand-alone	255 seconds	Without a cluster partner, the suggested value is always 255 seconds. The longer time-out value is to give the storage system time to recover from the failure.
nodev-tmo	Single Image, Partner, Dual Fabric	30 seconds	In Single-Image, Partner, or Dual Fabric mode, use a shorter value to enable a faster failover mode. Because the host has access to storage system interconnects, the host multipathing protocol starts using the secondary paths that use the interconnects. As a result, I/O can start faster in most cases. Exceptions: If you have an earlier version of the Host Utilities, you should use a longer timeout. ◆ For FC Solaris Attach Kit 2.0, use a value of 120 seconds. ◆ For FC Solaris Attach Kit 2.0 with the Support Matrix Update applied, use a value of 45 seconds.

cfmode-dependent values for parameters on systems using DMP with LPFC drivers			
Parameter	Cfmode	Value	Description
nodev-tmo	Standby, Mixed	120 seconds	These modes require a longer time-out in order to re-establish I/O because the storage system handles the LUN takeover. Having a longer time-out enables the host to wait while the storage system deals with the takeover. Exceptions: If you have an earlier version of the Host Utilities, you should use a longer time-out. ◆ For FC Solaris Attach Kit 2.0, use a value of 120 seconds. ◆ For FC Solaris Attach Kit 2.0 with the Support Matrix Update applied, use a value of 120 seconds.
nodev-tmo	Stand-alone	255 seconds	Without a cluster partner, the suggested value is always 255 seconds. The longer time-out value is to give the storage system time to recover from the failure.

basic_config sets values in lpfc.conf that are not related to cfmode

(Veritas DMP with LPFC) The `basic_config` command changes several parameters in the `/kernel/drv/lpfc.conf` file to improve the performance of the Host Utilities. These changes are made as “best fit” settings; in some cases, you should adjust these settings based on your environment. The following table lists parameters that are changed without regard to the `cfmode` being used.

Parameter changes for DMP with LPFC drivers that are not related to cfmode		
Parameter	Value	Description
<code>dqfull-throttle-up-time</code>	4 seconds	The <code>dqfull-throttle-up-time</code> parameter specifies how frequently the driver should increment the current queue depth. It works with the <code>dqfull-throttle-up-inc</code> parameter to restore the queue depth back to its original value. The recommended value of 4 seconds means that every 4 seconds the driver increments the queue depth by the amount specified in the <code>dqfull-throttle-up-inc</code> parameter. For more information on how these two parameters work together, see the <code>dqfull-throttle-up-inc</code> parameter.

Parameter changes for DMP with LPFC drivers that are not related to cfmode		
Parameter	Value	Description
dqfull-throttle-up-inc	2	The <code>dqfull-throttle-up-inc</code> parameter tells the driver by how much it should increment the current queue depth. It works with the <code>dqfull-throttle-up-time</code> parameter to restore the queue depth back to its original value. You use this parameter to increment the current queue depth if both of the following are true: ◆ The current queue depth does not equal the <code>lun-queue-depth</code> and ◆ The <code>driver stop_send_io</code> flag for that device is equal to 0. Note— Do not exceed the <code>lun-queue-depth</code> value for this parameter. If you use the Host Utilities recommended settings of 2 for the <code>dqfull-throttle-up-inc</code> parameter and 4 seconds for the <code>dqfull-throttle-up-time</code> parameter, the driver will increment the queue depth twice every four seconds. These values increase the queue depth, but at a slower rate than the default of one increment every second. The slower increase of the queue depth helps to allow the storage controller to recover from queue full conditions.

Parameter changes for DMP with LPFC drivers that are not related to cfmode		
Parameter	Value	Description
lun-queue-depth	16	This parameter specifies the number of outstanding commands for each FC LUN. The recommended value is 16; however, the size of your environment can affect the value you use here. You may need to tweak this setting based on your environment. The <code>lun-queue-depth</code> parameter works with the <code>tgt-queue-depth</code> parameter. The value of the <code>lun-queue-depth</code> parameter cannot exceed the value of the <code>tgt-queue-depth</code> command. For example, suppose a target has three LUNs. Each LUN has a queue depth of 16, The target queue depth is 32. If two of the LUNs fill their queues, resulting in a total of 32 queued commands, then no other commands can be queued up on any of the LUNs. The total LUN queue depth must fall below the target queue depth before additional commands can be added any of the LUN queues.
num-bufs	1024	This parameter specifies the number buffers to allocate. These buffers consume physical memory and are also used by the device driver to process loop initialization and rediscovery activities. The Host Utilities increase the value for this parameter from the default value to enable larger buffers for the Fibre Channel Extended Link services. The increased value also provides more buffer space for queuing FC commands.
num-iocbs	1024	This parameter specifies 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. Having too few IOCBs can temporarily prevent the driver from communicating with the adapter, which lowers performance. The Host Utilities increase the value for this parameter from the default value to ensure that an adequate number of control blocks are provided.

Parameter changes for DMP with LPFC drivers that are not related to cfmode		
Parameter	Value	Description
scan-down	0	The <code>basic_config</code> command resets the value to the driver default, which is 0. The <code>scan-down</code> parameter deals with loop topology, either public or private. It uses a static <code>AL_PA</code> map table. Setting the <code>scan-down</code> to 0 means that devices on a private loop are scanned from lowest to highest. The driver assigns the loop ID to the device by scanning the static <code>AL_PA</code> map table from low <code>AL_PA</code> to high <code>AL_PA</code>.
tgt-queue-depth	256	This parameter specifies the number of outstanding commands for each FC target. The recommended value is 256; however, the size of your environment can affect the value you use here. You may need to tweak this setting based on your environment. The <code>tgt-queue-depth</code> works with the <code>lun-queue-depth</code> parameter. For more information, see the <code>lun-queue-depth</code> parameter. Exceptions: For the Host Utilities 4.3 and earlier, the recommended value is 32.

Parameter changes for DMP with LPFC drivers that are not related to cfmode		
Parameter	Value	Description
topology		For the FC Solaris Host Utilities 4.3, the <code>basic_config</code> makes topology settings in the <code>/kernel/drv/lpfc.conf</code> file based on how the host is connected to the storage controller. When the host is connected to the storage controller through a fabric, the topology setting is forced to point-to-point. This ensures that the host will only try to connect to the storage controller as an <code>N_Port</code>. If the host is connected to the storage controller by way of a direct link (direct connect), the topology setting is set to force a loop mode connection. The port will only connect as a <code>NL_Port</code>. For earlier versions of the Host Utilities, the system administrator should decide which topology is best for initializing the Fibre Channel connection on your system. While you can use the <code>basic_config</code> command to set this value, by default, earlier versions of the Host Utilities do not modify this setting. To set this value using the <code>basic_config</code> command, enter the command: <pre>basic_config -topology = x</pre> Common values include: ◆ <code>topology=0</code> –(default value) attempt arbitrated loop mode; if that fails, attempt point-to-point mode ◆ <code>topology=2</code> –point-to-point mode (switch) ◆ <code>topology=4</code> –arbitrated loop mode Note— The FC Solaris Attach Kit 2.0 includes a wizard that asks you about your topology settings. Later kits let you use the <code>-topology</code> command line option.

**basic_config sets
values in
/etc/system**

(Veritas DMP with LPFC) The `basic_config` command makes the following changes to parameters in the `/etc/system` file.

Parameter changes for Veritas DMP with LPFC drivers	
Value	Description
<code>sd_io_time</code>	The <code>sd_io_time</code> variable specifies the time-out value for each I/O of <code>sd_io_time</code>. It is the limiter on how long an I/O can be outstanding before an error condition is returned. When a target device is no longer responding to commands, each I/O operation can take up to <code>sd_io_time * sd_retry_count</code> before the HBA driver times it out. Internal testing has shown that using the same value here as the one entered for the HBA driver's <code>nodev-tmo</code> parameter works well. This means that, in clustered environments, the value varies depending on your <code>cfmode</code>. In stand-alone environments, a longer value is required to allow more time for failover and giveback of the controllers. The current parameter settings are as follows: ◆ For Standby or Mixed <code>cfmode</code>, the setting is 120 seconds. To set this, enter: ❖ <code>set sd:sd_io_time=0x78</code> ◆ For Single-image, Partner, or Dual Fabric <code>cfmode</code>, the setting is 30 seconds. To set this, enter: ❖ <code>set sd_io_time=0x1e</code> ◆ For Stand-Alone systems, the setting is 255 seconds. To set this, enter: <code>set sd:sd_io_time=0xff</code>

Parameter changes for Veritas DMP with LPFC drivers	
Value	Description
	Exceptions for the FC Solaris Attach Kit 2.0: The following exceptions apply in clustered environments using the FC Solaris Attach Kit 2.0: ◆ For all cfmodes when you are using only the FC Solaris Attach Kit 2.0, enter a value of 120 seconds. ◆ For Standby or Mixed cfmode when you are using the FC Solaris Attach Kit 2.0 with the Support Matrix Update applied, enter a value of 120 seconds. ◆ For Single Image, Partner, or Dual Fabric cfmodes when you are using the FC Solaris Attach Kit 2.0 with the Support Matrix Update applied, enter a value of 45 seconds. Note_____ For information about what affects the <code>sd_io_time</code> value, see the description of the <code>nodev-tmo</code> parameter in “basic_config sets values in lpfc.conf based on cfmode” on page 3. _____
<code>sd_retry</code>	This variable specifies the number of retries that the SCSI driver will attempt before declaring a failure. Solaris 9 systems: This variable is changed on Solaris 9 systems only. Internal testing has shown that a value of 5 works well when you are running the Host Utilities on a Solaris 9 system. <pre>set sd:sd_retry_count=5</pre> Solaris 10 systems: The Host Utilities use the default value of 3 on Solaris 10 systems. There is no way to change this value on Solaris 10 systems. This value lives in the SCSI driver and does not show up in the file.

create_binding sets values in lpfc.conf

(Veritas DMP with LPFC) The `create_binding` program enables you to bind the WWPNs of a single storage system's active ports to a target ID and saves these bindings in the `/kernel/drv/lpfc.conf` file. Only one `fc-bind-WWPN` entry is allowed. That entry goes in the `lpfc.conf` file and sets a list of target portname bindings. This entry is based on recommendations from Emulex. The settings allow the LPFC driver to hand LUNs to the SCSI driver.

There also needs to be a corresponding entry in the `sd.conf` file to match the information added to the `lpfc.conf` file. Every LUN you create to use on the host must have a corresponding entry in the `/kernel/drv/sd.conf`. You must manually set that entry.

Currently, the `create_binding` program adds the line to the `lpfc.conf` file:

```
fcplib-bind-WWPN="50a980010300e073:lpfc0t0;"
```

which sets the following:

- ◆ World Wide Port Name (WWPN) is set to `50a980010300e073`
- ◆ The `lpfc` port instance is set to `lpfc0`.

Note

One driver controls all instances of the existing initiator.

- ◆ The target number per initiator on the host is set to `t0`. This tells the host to start at 0 and count up.

Sample entry in the `lpfc.conf` file: The following is a sample entry in the `lpfc.conf` file. The host gets the target numbers per initiator from this file.

```
fcplib-bind-WWPN="500a098a83d09b2d:lpfc0t0",  
                  "500a098183e0b481:lpfc0t1",  
                  "500a098383e0b481:lpfc1t0",  
                  "500a098c83d09b2d:lpfc1t1",  
                  "500a098a83e0b481:lpfc0t2",  
                  "500a098183d09b2d:lpfc0t3",  
                  "500a098383d09b2d:lpfc1t2",  
                  "500a098c83e0b481:lpfc1t3";
```

When the host is connected to multiple storage systems: If the host is connected to multiple storage systems, you must create an entry for each storage system and for each LUN mapped to the host in the `sd.conf` file.

Entry in the `sd.conf` file: If you want to enter the `lpfc.conf` information in the `/kernel/drv/sd.conf` file, the corresponding entry would use the format:

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

Note

None of the Host Utilities tools place this information in the `sd.conf` file. If you want it in the `sd.conf` file, you must manually edit that file.

Recommended settings for Veritas DMP with Sun native drivers

basic_config sets values in `ssd.conf` for DMP with Sun native drivers

(Veritas DMP with Sun native drivers) If you are using Veritas DMP with Sun native drivers, the Host Utilities require that you use single-image cfmode. The `basic_config` command makes the following changes to parameters in the `/kernel/drv/ssd.conf` file.

Note

The Solaris Host Utilities 4.3 introduced support for Veritas DMP with Sun native drivers. Prior to that release, this environment was not supported.

Parameter changes for DMP with Sun native drivers		
Parameter	Value	Description
<code>busy_retries</code>	30	Setting this value to 30 assists the host in surviving a storage controller fail-over.
<code>not_ready_retries</code>	300	Setting this value to 300 gives the Sun Cluster 3.1 time to allow for a storage controller transition. Testing has shown that hosts using Sun Cluster were unable to survive a storage controller fail-over without increasing this variable to 300.
<code>reset_retries</code>	30	Setting this value to 30 assists the host in surviving a storage controller fail-over without having to increase this variable.
<code>throttle_max</code>	8	Setting this value to 8 prevents issues seen with queue full conditions.
<code>throttle_min</code>	2	This value needs to be lower than the <code>throttle_max</code> value. Since the <code>throttle_max</code> value must be set to 8, this value was lowered to 2. The default for this value is 8.

Recommended Host Utilities settings for MPxIO environment

basic_config sets values in ssd.conf, sd.conf

(MPxIO only) If you are using MPxIO, the Host Utilities require that you use single-image cfmode. The `basic_config` command makes the following changes to parameters in either `/kernel/drv/ssd.conf` file (SPARC processors) or `/kernel/drv/sd` (x86/64 processors) depending on the options you use with it.

Parameters changed for MPxIO environments		
Parameter	Value	Description
<code>busy_retries</code>	30	Setting this value to 30 assists the host in surviving a storage controller fail-over without having to increase this variable.
<code>not_ready_retries</code>	300	Setting this value to 300 gives the Sun Cluster 3.1 time to allow for a storage controller transition. Testing has shown that hosts using Sun Cluster were unable to survive a storage controller fail-over without increasing this variable to 300.
<code>reset_retries</code>	30	Setting this value to 30 assists the host in surviving a storage controller fail-over without having to increase this variable.
<code>throttle_max</code>	64	This is the value recommended by Sun Microsystems for SAN Foundation Suite tunables.
<code>throttle_min</code>	8	This is the default for this value for Solaris 10 Update 3. NetApp recommends keeping it at 8 to maintain backward compatibility with Solaris 10 Update 2.

Recommended settings when using Veritas VxVM and DMP

Recommendations for Veritas values

You can run the Host Utilities with Veritas Volume Manager (VxVM) and Dynamic Multipathing (DMP). NetApp recommends the following values:

- ◆ The restore daemon settings. The values for the DMP restore daemon's policy and interval tuneables determine how frequently the daemon checks paths between the host and the storage system and which paths it checks. You need to supply values for the following two settings:
 - ❖ The restore daemon interval. Specify the number of seconds between polls.

If you are using a version of the Veritas Storage Foundation prior to 5.0 (and Host Utilities 4.2), you get a faster recovery by making the restore daemon poll more often. Systems set to a longer polling interval take longer to detect path recoveries, which can impact storage system performance during a failover operation.
 - ❖ The restore daemon policy. Specify which paths the restore daemon should check during polls.

If you set the value to `check_disabled`, the daemon only checks the disabled paths. If you set this to `check_all`, it checks all the paths between the storage system and the host.

With Veritas Storage Foundation 5.0, `check_disabled` provides better performance. If you have a version of the Veritas Storage Foundation prior to 5.0, `check_all` results in faster recovery times during failovers and failbacks.

Note

The values NetApp recommends for the restore daemon are based on testing done at the time the Host Utilities or Attach Kit was created.

- ◆ All other settings. Use the default values.

Restore daemon settings (Host Utilities 4.2 and greater): Versions 4.2 and higher of the FC Solaris Host Utilities for Veritas use Veritas Storage Foundation 5.0, which fixes an issue that occurred in earlier Storage Foundation versions. For Host Utilities using Storage Foundation 5.0, NetApp recommends the following values for the restore daemon:

- ◆ restore daemon interval: 60
- ◆ restore daemon policy: `check_disabled`

Restore daemon settings (Host Utilities 4.0/ 4.1, Attach Kit 3.0): For the FC Solaris Host Utilities 4.0/4.1 for Veritas and for the FC Solaris Attach Kit 3.0, NetApp recommends the following values for the restore daemon:

- ◆ restore daemon interval: 60
- ◆ restore daemon policy: check_all

For example, when you enter `vxddmpadm stat restore`, you should see the following:

```
The number of daemons running : 1
The interval of daemon: 60
The policy of daemon: check_all
```

For information on how to set these values, see the Veritas documentation.

Exceptions for the FC Solaris Attach Kit 2.0: If you have an FC Solaris Attach Kit 2.0, use the following settings:

- ◆ restore daemon interval: 300
- ◆ restore daemon policy: check_disabled (default)

For example, when you enter `vxddmpadm stat restore`, you should see the following:

```
The number of daemons running : 1
The interval of daemon: 300
The policy of daemon: check_disabled
```

If you have an FC Solaris Attach Kit 2.0 with the Support Matrix Update enabled, use these settings:

- ◆ restore daemon interval: 60
- ◆ restore daemon policy: check_disabled