

Host Settings Affected by Solaris Host Utilities (Veritas)

FCP Solaris Host Utilities for Veritas changes to host settings

In some cases, the FCP Solaris Host Utilities for Veritas recommend that settings on the host be modified. These recommendations are based on research and internal testing done at NetApp. 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). Also, the recommendations may vary depending on whether you are running the current, 4.x version of the Host Utilities or an early version of the Attach Kit, such as FCP Solaris Attach Kit 2.0.

For Veritas, the Host Utilities recommend that you use the values provided in the Veritas Release Notes except for the restore daemon settings. See [“Recommendations for Veritas settings”](#) on page 8 for the values you should use for the restore daemon's policy and interval tuneables.

This document provides a summary of the suggested host and Veritas settings.

Tools to modify host settings for FCP Solaris Host Utilities for Veritas

The FCP Solaris Host Utilities for Veritas provide two tools that modify settings on the host:

Command	Description
<code>basic_config</code>	Modifies parameters in the following files: ◆ <code>/kernel/drv/lpfc.conf</code> ◆ <code>/etc/system</code> Note_____ You can manually set these parameters instead of using <code>basic_config</code>, if you choose.
<code>create_binding.pl</code>	Modifies parameters in <code>/kernel/drv/lpfc.conf</code> to create persistent bindings. Note_____ You can use other tools, such as <code>lputil</code> or <code>HBAnyware</code>, to create persistent bindings.

**(Clustered)
basic_config sets
values in lpfc.conf**

On clustered systems the basic_config command makes the following changes to parameters in the /kernel/drv/lpfc.conf file based on the cfmodes for the system:

Clustered Systems			
Parameter	Cfmodes	Value	Description
dqfull-throttle-up-time	All modes	4 seconds	The dqfull-throttle-up-time parameter tells the driver when it should increment the current queue depth. Internal testing has shown that a value of 4 seconds works well for this parameter.
dqfull-throttle-up-inc	All modes	2 seconds	The dqfull-throttle-up-inc parameter tells the driver by how much it should increment the current queue depth. Internal testing has shown that a value of 2 works well for this parameter.
linkdown-tmo	All modes	30 seconds	The basic_config 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.
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 right away. Exceptions: If you have an earlier version of the Host Utilities, you should use a longer timeout. ◆ For FCP Solaris Attach Kit 2.0, use a value of 120 seconds. ◆ For FCP Solaris Attach Kit 2.0 with the Support Matrix Update applied, use a value of 45 seconds.

Clustered Systems			
Parameter	Cfmode	Value	Description
nodev-tmo	Standby, Mixed	120 seconds	These modes require a longer timeout in order to re-establish I/O because the storage system handles the LUN takeover. Having a longer timeout 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 timeout. ◆ For FCP Solaris Attach Kit 2.0, use a value of 120 seconds. ◆ For FCP Solaris Attach Kit 2.0 with the Support Matrix Update applied, use a value of 120 seconds.
lun-queue-depth	All modes	16	This parameter specifies the number of outstanding commands for each FCP LUN. Internal testing has shown that a value of 16 works well for this parameter.
num-bufs	All modes	1024	This parameter specifies the number buffers to allocate. Internal testing has shown that a value of 1024 works well for this parameter.
num-iocbs	All modes	1024	This parameter specifies the number of IOCB (Input/Output command blocks) buffers to allocate. Internal testing has shown that a value of 1024 works well for this parameter.
scan-down	All modes	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 tells the driver to assign 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>.

Clustered Systems			
Parameter	Cfmode	Value	Description
tgt-queue-depth	All modes	32	This parameter specifies the number of outstanding commands for each FCP target. Internal testing has shown that a value of 32 works well for this parameter.
topology	All modes		The system administrator should decide which topology is best for initializing the Fibre Channel connection on your system. By default, the Host Utilities do not modify this setting. You can use the <code>basic_config</code> command to set this value, if you wish. Enter the command: <pre>basic_config -topology = x</pre> Common values include: ◆ topology=0 –(default value) attempt arbitrated loop mode; if that fails, attempt point-to-point mode ◆ topology=2 –point-to-point mode (switch) ◆ topology=4 –arbitrated loop mode Exceptions: If you have an earlier version of the Host Utilities, you should use a longer timeout. ◆ For FCP Solaris Attach Kit 2.0, use a value of 120 seconds. Note _____ The FCP 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. _____

**(Stand-alone)
basic_config sets
values in lpfs.conf**

On a stand-alone system, the `basic_config` command makes the following changes to parameters in the `/kernel/drv/lpfc.conf` file:

Stand-Alone Systems		
Parameter	Value	Description
linkdown-tmo	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	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.

The other parameters that apply to a stand-alone system should have the same values as those listed for a clustered system.

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

The `basic_config` command makes the following changes to parameters in the `/etc/system` file.

Value	Description
sd_io_time	The <code>sd_io_time</code> variable specifies the time-out value for disk operations. It is the limiter on how long an I/O can be outstanding before an error condition is returned. 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 <code>nodev-tmo</code> 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>

Value	Description
	◆ For Stand-Alone systems, the setting is 255 seconds. To set this, enter: <pre>set sd:sd_io_time=0xff</pre> Exceptions for the FCP Solaris Attach Kit 2.0: The following exceptions apply in clustered environments using the FCP Solaris Attach Kit 2.0: ◆ For all cfmodes when you are using only the FCP Solaris Attach Kit 2.0, enter a value of 120 seconds. ◆ For Standby or Mixed cfmode when you are using the FCP 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 FCP Solaris Attach Kit 2.0 with the Support Matrix Update applied, enter a value of 45 seconds. Note_____ For information about what affects the sd_io_time value, see the description of the nodev-tmo parameter in “(Clustered) basic_config sets values in lpfc.conf” on page 2 and “(Stand-alone) basic_config sets values in lpfs.conf” on page 5.
sd_retry	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 on Solaris 10 systems. This value lives in the SCSI driver and does not show up in the file.

create_binding.pl sets values in lpfc.conf

The create_binding.pl script 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 fcp-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.pl` script adds the line to the `lpfc.conf` file:

```
fcplib=lpfc:lpfc0
```

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=lpfc:lpfc0: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.

Recommendations for Veritas settings

NetApp 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 givebacks.

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 values recommended in the Veritas Release Notes.

Restore daemon settings (Host Utilities 4.2): The FCP Solaris Host Utilities 4.2 for Veritas use Veritas Storage Foundation 5.0, which fixes an issue that occurred in earlier Storage Foundation versions. For Host Utilities 4.2, 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 FCP Solaris Host Utilities 4.0/4.1 for Veritas and for the FCP 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 `vxdmpadm 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. Instructions are also included in the FCP Solaris Host Utilities 4.2 for Veritas Installation and Administration Guide.

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

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

For example, when you enter `vxdmpadm 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 FCP Solaris Attach Kit 2.0 with the Support Matrix Update enabled, use these settings:

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